CHALMERS
UNIVERSITY OF TECHNOLOGY

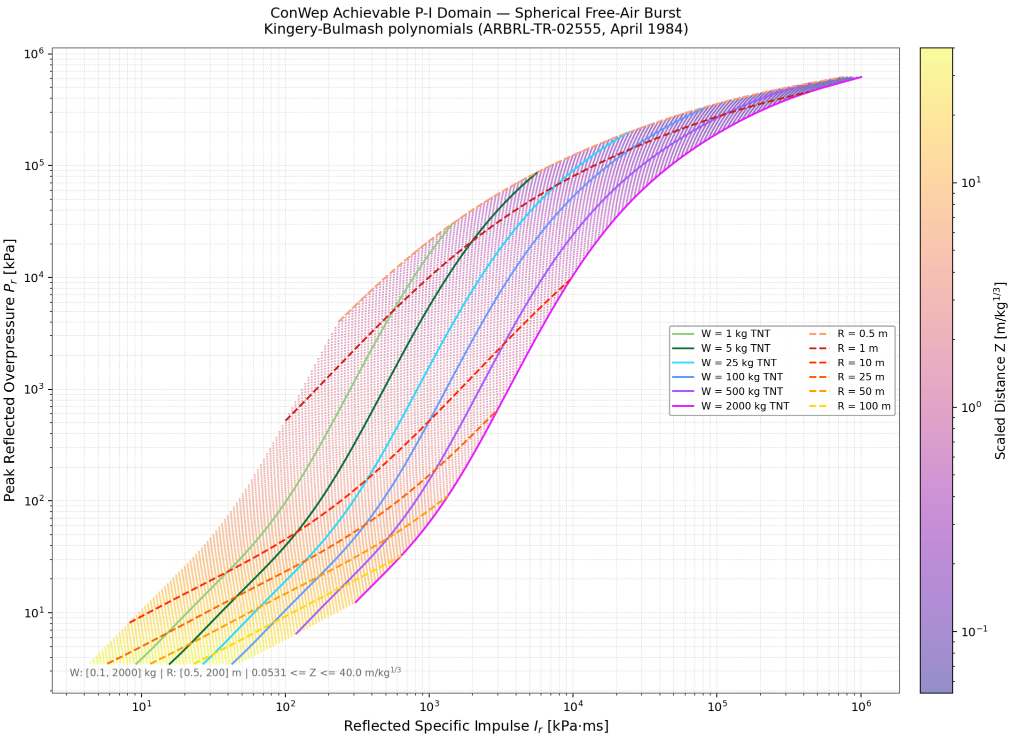

Graphics-assisted tool for the automated numerical modelling of blast-loaded masonry structures

An automated finite-element workflow for constructing pressure–impulse failure diagrams of unreinforced masonry walls

Master's thesis in Structural Engineering and Building Technology

Ahmed Ala-Abadi
Jens Le

DEPARTMENT OF ARCHITECTURE AND CIVIL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

MASTER'S THESIS ACEX30

Graphics-assisted tool for the automated numerical modelling of blast-loaded masonry structures

An automated finite-element workflow for constructing pressure–impulse failure diagrams of unreinforced masonry walls

Master's Thesis in the Master's Programme Structural Engineering and Building Technology

Ahmed Ala-Abadi
Jens Le

Department of Architecture and Civil Engineering
Division of Structural Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026

Graphics-assisted tool for the automated numerical modelling of blast-loaded masonry structures
An automated finite-element workflow for constructing pressure–impulse failure diagrams of unreinforced masonry walls

Master's Thesis in the Master's Programme Structural Engineering and Building Technology

Ahmed Ala-Abadi
Jens Le

Supervisor: Michele Godio, RISE Research Institutes of Sweden
Examiner: Joosef Leppänen, Department of Architecture and Civil Engineering

Examensarbete ACEX30
Institutionen för Arkitektur och Samhällsbyggnadsteknik
Chalmers Tekniska Högskola, 2026

Department of Architecture and Civil Engineering
Chalmers University of Technology
SE-412 96 Gothenburg
Sweden
Telephone +46 31 772 1000

Cover: ConWep-achievable pressure–impulse domain for a spherical free-air burst, computed from the Kingery–Bulmash polynomials (see Section 3.2.2).

Typeset in LaTeX
Gothenburg, Sweden 2026

Graphics-assisted tool for the automated numerical modelling of blast-loaded masonry structures
An automated finite-element workflow for constructing pressure–impulse failure diagrams of unreinforced masonry walls

Master's Thesis in the Master's Programme Structural Engineering and Building Technology

Ahmed Ala-Abadi

Jens Le

Department of Architecture and Civil Engineering
Division of Structural Engineering
Chalmers University of Technology

ABSTRACT

Pressure-impulse diagrams provide a compact representation of the failure boundary of a structural element under blast loading, but their construction for masonry walls is expensive, where each point requires a full dynamic finite-element analysis. This thesis develops a methodology and an accompanying graphical tool that automates the construction of pressure-impulse failure boundaries for unreinforced masonry walls. The objective is to establish a reproducible workflow for arbitrary wall configurations.

The methodology selects components along three axes. Masonry is represented by a simplified-micro model with bricks modelled by Concrete Damaged Plasticity with rate-dependent hardening, along with mortar joints as cohesive interfaces. The blast load is implemented in two modes: validation mode applies a ConWep load for scenarios with known charge weight and standoff, and sweep mode applies a single-coefficient exponential pulse parameterised directly by the target peak pressure and impulse. The sweep algorithm is an adaptive Bayesian boundary tracer that combines strategic seeding, fix-pressure and fix-impulse bisection, and parametric hyperbolic fitting.

The methodology is demonstrated on an arbitrary single-wythe unreinforced masonry wall of 1200 × 1155 × 50 mm, with zero-thickness mortar joints, in-plane fixed and out-of-plane pinned supports, and an axial pre-stress. The criterion used for failure is taken at the UFC 3-340-02 transition between the reusable and non-reusable damage states. The sweep identifies nine boundary points in sixty finite-element evaluations across approximately 35 hours of wall-clock time, spanning the impulsive, dynamic, and quasi-static regimes. Fitted hyperbolic curves bracket the pressure asymptote between 1.33 and 3.05 kPa depending on parametric form, with the impulse asymptote near 55 kPa·ms. The graphical tool achieves its design goal of making the modelling process accessible without case-by-case manual setup.

Key words:
masonry walls, blast loading, pressure-impulse diagram, adaptive sampling, Bayesian inference, finite element analysis, automation

Contents

Acknowledgements

This thesis was carried out primarily at RISE Research Institutes of Sweden, within the project “Blast capacity of masonry walls” under the area “CFORT D1 – Explosion and weaponry effects”. We gratefully acknowledge this project for funding the work, the supervision, and our compensation, without which the thesis would not have been possible.

We would like to thank our supervisor at RISE, Michele Godio, for proposing the project and for his guidance, technical insight, and continued support throughout every stage of the work. We also thank our examiner at Chalmers University of Technology, Joosef Leppänen, for his oversight and feedback during the project.

We are grateful to Viktor Peterson at FOI, who as part of the supervision panel contributed actively to our meetings and offered valuable input that helped shape the direction of the work.

Ahmed Ala-Abadi & Jens Le, Gothenburg, June 2026

List of Acronyms

Below is the list of acronyms that have been used throughout this thesis listed in alphabetical order:

ALE	Arbitrary Lagrangian–Eulerian
AK-MCS	Active-learning reliability method combining Kriging and Monte Carlo Simulation
BIC	Bayesian Information Criterion
CDP	Concrete Damaged Plasticity
CEL	Coupled Eulerian–Lagrangian
ConWep	Conventional Weapons Effects program
DIF	Dynamic Increase Factor
DOF	Degree of Freedom
EGRA	Efficient Global Reliability Analysis
FE	Finite Element
FEM	Finite Element Method
GUI	Graphical User Interface
JWL	Jones–Wilkins–Lee (equation of state)
MCMC	Markov-Chain Monte Carlo
P–I	Pressure–Impulse
RBSM	Rigid Body-Spring Model
RSS	Residual Sum of Squares
SDOF	Single-Degree-of-Freedom
TNT	Trinitrotoluene
UFC	Unified Facilities Criteria
UHPFRC	Ultra-High Performance Fibre Reinforced Concrete
URM	Unreinforced Masonry

Nomenclature

Below is the nomenclature of the symbols used throughout this thesis, grouped by topic.

Blast Loading

W	Equivalent TNT charge mass
R	Standoff distance
Z	Scaled distance, $Z = R/W^{1/3}$
P_{so}	Peak incident overpressure
P_r	Peak reflected overpressure
P_{peak}	Peak pressure of the blast pulse
t_d	Positive-phase duration
t_a	Shock arrival time
i_s	Incident specific impulse
I_r	Reflected specific impulse
b	Dimensionless waveform decay coefficient (modified Friedlander equation)
κ	Decay coefficient of the single-parameter exponential pulse, $\kappa = P_{peak}/I$
C_n	Fitted Kingery–Bulmash polynomial coefficients
$K_0,\ K_1$	Shift and scale constants of the Kingery–Bulmash polynomial
θ	Angle of blast incidence (Section 2.4.2)

Dynamic Finite-Element Analysis

$\boldsymbol{M}$	Global mass matrix
$\boldsymbol{C}$	Damping matrix
$\boldsymbol{K}$	Global stiffness matrix
$\boldsymbol{F}(t)$	Time-dependent external load vector
$\boldsymbol{u}$	Nodal displacement vector
$\dot{\boldsymbol{u}}$	Nodal velocity vector
$\ddot{\boldsymbol{u}}$	Nodal acceleration vector

Masonry Wall and Material Properties

H	Wall height (clear span of the one-way wall)
L	Wall length
t	Wall thickness
H/t	Slenderness ratio
h_{brick}	Brick unit height
l_{full}, l_{edge}	Full-brick and edge-brick length
n_{layers}	Number of brick courses
n_{full}	Number of full bricks per course
g_{mortar}	Mortar joint gap (zero in the simplified micro-model)
$t_{n,0}$	Mortar tensile bond strength
G_f	Mortar fracture energy
ε	Strain
$\dot{\varepsilon}$	Strain rate

Damage Assessment

θ	Chord (support) rotation
θ_{mid}	Peak midspan support rotation
θ_{fail1}	Non-reusable rotation threshold ($1.0°$)
θ_{fail2}	Reusable rotation threshold ($0.5°$)
δ	Out-of-plane midspan displacement
δ_{mid}	Peak midspan deflection
$\delta_{collapse}$	Collapse deflection threshold (equal to wall thickness)
$G(X)$	Limit-state (damage) function

Pressure–Impulse Diagram and Adaptive Sweep

P	Peak pressure (P–I diagram ordinate)
I	Total specific impulse (P–I diagram abscissa)
$P(I)$	Fitted failure-boundary curve
a	Pressure asymptote of the P–I curve (quasi-static capacity)
c	Impulse asymptote of the P–I curve (impulsive capacity)
b	Curvature coefficient of the hyperbolic fit
n	Curvature exponent of the four-parameter hyperbola
P_0, I_0	Pressure and impulse asymptotes (Chapter 6 notation; $P_0 \equiv a$, $I_0 \equiv c$)

k	Number of model parameters
n_{pts}	Number of located boundary points
ΔBIC	BIC difference between candidate models
τ	Integrated autocorrelation time of the MCMC chain
n_{steps}	MCMC chain length
$\alpha_{var}(I)$	Predictive-uncertainty acquisition score
$\alpha_{cov}(I)$	Coverage acquisition score
$\alpha(I)$	Blended acquisition score
w	Coverage-bias weight
$\sigma_{\log_{10} P}$	Predictive standard deviation of the fit in $\log_{10} P$ space

1 Introduction

1.1 Background

Explosions load buildings with a regime conventional design does not anticipate: a pressure pulse of extreme magnitude and very short duration, arriving without warning. Such events arise from accidental sources, such as gas explosions and transport incidents, as well as from deliberate ones. Assessing the vulnerability of structures that were never designed for them has become an established concern of structural engineering, and unreinforced masonry (URM) is central to it. URM is one of the most widespread materials in the existing building stock and among the most vulnerable to blast: it has little tensile strength, brittle mortar joints and limited ductility, so a wall struck on its face is loaded out-of-plane in the direction in which it is weakest. Its out-of-plane failure, with the projection of debris into the space behind it, is a recognised dominant cause of injury in explosion events.

The masonry wall is, however, a difficult object to analyse. It is a heterogeneous assembly of stiff brick units and weak mortar joints, and its response is governed by the joints, where cracking, sliding and debonding concentrate long before the units themselves fail. The boundary conditions are rarely simple, and the blast load itself must be characterised in magnitude, distribution and time history before any predictive analysis can begin. As a result, modelling a single scenario consumes substantial analyst time before any result is produced.

A single scenario, moreover, is seldom the question protective assessment poses. The relevant question is how a wall's survival depends on the loading across a range of credible threats, and the pressure–impulse (P–I) diagram is the established means of expressing that answer. It is a single curve in the plane of peak pressure and impulse, separating the load combinations an element survives from those it does not, and resolving naturally into impulsive, dynamic and quasi-static regimes.

Conventionally, P–I curves are produced not by finite-element analysis but by simplified hand calculation, using single-degree-of-freedom idealisations of the kind codified in the Swedish design rules [41]. These are inexpensive but restricted to simple resistance functions and idealised pulses, and are difficult to apply to the joint-dominated out-of-plane response of masonry.

Resolving the masonry response at the unit–mortar level instead requires finite-element analysis, and there the difficulty is cost: each point on the boundary is the outcome of a full non-linear dynamic analysis, and a boundary resolved on a dense grid may require hundreds of them. When each analysis is a finite-element model of a heterogeneous masonry wall built and post-processed by hand, the cost of a single diagram becomes a strong disincentive to the parametric studies protective assessment requires.

This barrier is reflected in the literature. Work on P–I diagrams is dominated by reinforced-concrete and steel members, while masonry-specific studies are scarce, and those that exist rely either on homogenised macro-models that cannot resolve the unit–mortar interface failures central to URM cracking, or on dense, fixed sampling schemes. No established, reproducible workflow constructs masonry P–I failure diagrams without case-by-case manual finite-element setup. Closing that gap is the concern of this thesis.

1.2 Previous research

This thesis is conducted within the RISE Research Institutes of Sweden project *Blast capacity of masonry walls*, falling under the research area CFORT D1 – Explosion and weaponry effects, which was established to increase knowledge on the blast response of unreinforced masonry structures in Sweden. Prior work within the project has covered free far-field blast tests [21], low-velocity and moderate-velocity out-of-plane impact tests [17, 19], four-point bending tests under varying support conditions [20], and the development of a validated finite element meso-scale modelling approach [18]. The present thesis extends this programme by constructing an automated finite element pipeline for generating pressure-impulse diagrams of URM walls, providing a blast vulnerability assessment tool not previously realised within the project.

1.3 Purpose and aim

The purpose of this thesis is to lower the analyst-time barrier to performing parametric blast assessments of unreinforced masonry walls, by developing an automated modelling workflow built on top of a commercial finite-element solver. The intent is to make building, running and post-processing a blast-loaded masonry analysis sufficiently inexpensive that families of scenarios, represented as pressure–impulse (P–I) diagrams, can be constructed without case-by-case manual setup, and to establish a reproducible procedure applicable to arbitrary wall configurations rather than to validate any one configuration against experiment.
The aim of the project is to develop a schematisation and parameterisation tool that assists analysts in constructing reliable numerical models of blast-loaded masonry walls in an automated way. The tool is required to generate a non-linear finite-element model of a masonry wall in ABAQUS from a parametric description of its geometry, bond pattern and boundary conditions; to represent the blast load both for the reproduction of real explosion events and for parametric exploration of the load plane; and to drive an adaptive procedure that traces a P–I failure boundary from the resulting analyses. It is exercised by producing a pressure–impulse failure diagram for a representative wall configuration.

1.4 Limitations

To bound the complexity of the implementation, the developed tool is restricted to blast loading, excluding impact scenarios and environmental actions such as temperature variation. The structural scope is limited to a single masonry wall; adjacent transverse walls, floors and surrounding structures are removed or simplified into boundary conditions on the wall surface. The model further assumes idealised material and geometric properties, neglecting construction defects and material variability. Finally, the work establishes and demonstrates a methodology: the failure boundary produced for the representative wall is checked for physical plausibility and internal consistency but is not validated against a physical test.

1.5 Research Novelty

The contribution of this thesis is methodological, and does not reside in any single algorithmic component. Bisection root-finding, the Bayesian Information Criterion, Markov-chain Monte Carlo and adaptive Bayesian sampling are all long-established tools. The novelty lies in their composition into a single workflow, in its application to a problem for which it has not previously been used, and in the automation that makes it usable.

Three aspects may be distinguished. First, the thesis delivers an end-to-end automated route from a parametric description of a masonry wall to a traced P–I failure boundary. P–I diagrams for masonry have previously been constructed as bespoke, case-by-case analyses; here the construction is reduced to a configurable and reproducible procedure, operated through a graphical interface, that requires no manual finite-element setup. Second, it applies an adaptive, uncertainty-driven boundary-tracing strategy to a masonry P–I diagram: a parametric curve is fitted to the points found so far, and the next evaluation is placed where that fit is least certain. The strategy is distinctive in generating candidate evaluations along the current boundary estimate rather than across the whole load plane, and has not previously been applied to a masonry P–I boundary. Third, it adopts a dual-mode treatment of the blast load: physically grounded for the reproduction of real events, and a single-parameter exponential pulse for the sweep, whose peak pressure and impulse are independent by construction. This uncoupling allows the sweep to reach the quasi-static region of the (P,I) plane that fixes the pressure asymptote, which an empirical TNT-equivalent load model cannot access. Fourth, the methodology is made accessible through a graphical user interface integrated with ABAQUS/CAE, which collects wall geometry, bond pattern, boundary conditions, and sweep parameters through structured input dialogs and drives the complete pipeline without requiring the user to edit source code or construct finite-element input files directly. This lowers the analyst-time barrier to parametric blast assessment and makes the workflow reproducible across arbitrary wall configurations.

Taken together, these contributions place the micro-modelled finite-element construction of a masonry P–I diagram within an automated framework for the first time.

1.6 Organization of the thesis

Chapter 2 reviews the underpinning theory: dynamic finite-element analysis, damage assessment of blast-loaded masonry, the numerical modelling strategies for masonry, the representation of blast loads, and the construction of P–I diagrams. Chapter 3 sets out the methodology, selecting a masonry model, a blast-load representation and a P–I sweep algorithm, and documenting the modelling trials behind those selections. Chapter 4 describes the framework that integrates these components into the final automated tool, including the graphical user interface. Chapter 5 demonstrates the framework end-to-end through a case study on a representative single-wythe URM wall and reports the resulting P–I diagram. Chapter 6 discusses those results, Chapter 7 draws the conclusions, and Chapter 8 outlines directions for future work.

1.7 Societal, ethical and ecological aspects

Unreinforced masonry constitutes a substantial portion of the existing built environment across Europe and beyond, encompassing residential housing, schools, hospitals, and cultural heritage structures. Many of these buildings were designed without any consideration of blast or extreme dynamic loading, leaving large segments of the civilian population vulnerable in the event of accidental explosions, terrorist incidents, or armed conflict. A deeper understanding of how URM walls respond to blast pressure, and the conditions under which they transition from repairable damage to collapse, directly informs protective engineering decisions that can save lives. The pressure-impulse diagram framework developed in this thesis provides practitioners with a computationally efficient tool for rapid vulnerability assessment, supporting risk-informed decisions in both the design of new structures and the retrofitting of existing ones.

Beyond individual buildings, the resilience of urban infrastructure depends on the collective performance of its structural elements under extreme loading. Governments and civil defence organisations increasingly recognise that protecting critical infrastructure from blast threats is not solely a military concern but a broader societal imperative. By providing quantitative damage thresholds, this work contributes to an effort to embed blast resilience into civilian planning frameworks, enabling authorities to identify and prioritise vulnerable structures in densely populated areas.

Research in blast engineering carries dual-use implications: the same knowledge that enables the protection of civilians can, in principle, inform the design or deployment of explosive devices. This places an ethical responsibility on researchers to frame their contributions explicitly within a protective context. The present work is oriented entirely toward passive structural protection, understanding failure mechanisms and quantifying resistance, rather than toward the optimisation of threat parameters.

Computational simulation, as employed in this thesis, represents an ecologically preferable alternative to physical blast testing, which involves material consumption and environmental disturbance. By replacing experimental campaigns with finite element models, the direct environmental footprint of the investigation is minimised.

2 Theory

2.1 Dynamic Finite Element Analysis

2.1.1 The Equation of Motion

The fundamental objective of any dynamic finite element analysis is to solve the governing equation of motion for the system, which is mathematically expressed as:

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{C}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{F}(t) \tag{2.1}$$

where $\mathbf{M}$ is the global mass matrix, $\mathbf{C}$ is the damping matrix, $\mathbf{K}$ is the global stiffness matrix, and $\mathbf{F}(t)$ represents the time-dependent external load vector (e.g., the applied blast pressure). The terms $\ddot{\mathbf{u}}$, $\dot{\mathbf{u}}$, and $\mathbf{u}$ represent the nodal acceleration, velocity, and displacement vectors, respectively. The primary difference between numerical solvers lies in how they integrate this equation over time.

2.1.2 Implicit Solvers

Implicit solvers solve the equation of motion by iteratively determining the unknown displacements ($\mathbf{u}$) at the end of each time step. A key characteristic of this method is that it requires the system to reach global static equilibrium before advancing to the next increment. To achieve this, the global stiffness matrix ($\mathbf{K}$) must be inverted during each iteration. For models with large global stiffness matrices, this inversion requires extensive computational memory and time.

While implicit solvers are unconditionally stable and typically allow for very large time steps, this advantage is negated in blast-loaded structures. Because a blast event is nearly instantaneous, extremely small time increments are required to accurately capture the stress wave propagation. Forcing an implicit solver to invert a massive $\mathbf{K}$-matrix at these microsecond increments becomes computationally prohibitive. Furthermore, a critical limitation arises regarding the quasi-brittle nature of masonry. When the bricks or mortar lose load-bearing capacity under the blast pressure, their local stiffness drops to zero. This results in the solver attempting to invert a $\mathbf{K}$-matrix containing zeros, leading to a mathematical singularity where the equation cannot converge or be solved.

2.1.3 Explicit Solvers

In contrast to the implicit method, explicit solvers determine the future kinematic state of the system based solely on its current, known state. Instead of solving for displacements iteratively, the solver calculates nodal accelerations directly from the equation of motion. By utilizing a lumped (diagonal) mass matrix, the calculation is reduced to dividing the net internal and external forces by the mass ($\ddot{\mathbf{u}} = \mathbf{M}^{-1}(\mathbf{F} - \mathbf{C}\dot{\mathbf{u}} - \mathbf{K}\mathbf{u})$, or conceptually $a = F/m$). Because this method bypasses the inversion of the global stiffness matrix, it does not iterate to enforce global equilibrium.

The primary limitation of the explicit solver is that it is conditionally stable; it requires exceptionally tiny time increments to accurately capture the state of the system. If the time increment is too large, the solver will fail to capture the high-frequency physics of the impact, and the mathematical solution will explode. However, because the mathematics required to solve for acceleration per step (simple algebraic division) are drastically simpler than repeatedly inverting a global $\mathbf{K}$-matrix, the explicit solver is vastly more computationally efficient per time step. This makes it highly capable of handling the severe non-linearities, contact changes, and zero-stiffness fragmentation inherent to blast-loaded masonry. Which is why the explicit solver will be used for this thesis.

2.2 Damage Assessment of Blast-Loaded Masonry Walls

2.2.1 Damage Metrics and Performance Levels

There is currently no unified international standard for quantifying blast-induced damage in unreinforced masonry (URM) walls. The Eurocodes provide no explicit provisions for external blast events, and existing national standards vary significantly in scope and methodology.

The most widely adopted framework for damage assessment of blast-loaded masonry is the United States *Unified Facilities Criteria* UFC 3-340-02 [43], which defines damage severity through *support rotation* thresholds derived from the out-of-plane displacement response. Table 6-2 of the UFC specifies two performance levels for masonry walls: a *reusable* limit of $\theta \leq 0.5^\circ$ (for both one-way and two-way spanning walls), and a *non-reusable* limit of $\theta \leq 1.0^\circ$ for one-way spanning walls and $\theta \leq 2.0^\circ$ for two-way spanning walls. The UFC does not define an explicit collapse threshold beyond the non-reusable limit.

To extend this two-tier framework into a three-tier classification suitable for performance-based assessment, Wei, Huang, and Li [44] supplemented the UFC rotation criteria with a deflection-based collapse criterion. Collapse is defined as the state in which the peak midspan out-of-plane deflection equals the wall thickness, on the physical basis that arching equilibrium is lost once the lateral displacement exceeds the wall's own thickness. Wei et al.'s three damage states are defined through the limit-state functions:

$$
\begin{aligned}
G(X) &= \delta_{\text{collapse}} - \delta_{\text{predicted}} && \text{(collapse, major damage)} \\
G(X) &= \theta_{\text{fail1}} - \theta_{\text{predicted}} && \text{(non-reusable, intermediate damage)} \\
G(X) &= \theta_{\text{fail2}} - \theta_{\text{predicted}} && \text{(reusable, minor damage)}
\end{aligned}
\tag{2.2}
$$

where δ_{collapse} equals the wall thickness, $\theta_{\text{fail1}} = 1.0^\circ$ and $\theta_{\text{fail2}} = 0.5^\circ$ are the non-reusable and reusable rotation thresholds respectively for one-way spanning walls (per Wei et al.'s Table 2, adapted from UFC Table 6-2), and damage occurs when $G(X) \leq 0$. The chord rotation θ is computed geometrically as:

$$
\theta = \arctan\left(\frac{\delta}{H/2}\right) \tag{2.3}
$$

where δ is the peak out-of-plane midspan displacement and H is the clear span (wall height) between supports. This formulation derives from Table 3-5 of UFC 3-340-02, which idealises the deflected shape of one-way spanning elements as two rigid half-spans rotating about plastic hinges at the supports and at midspan.

The rotation thresholds in this framework are performance-based criteria for the wall itself; they apply regardless of the underlying resistance mechanism. For unreinforced masonry walls, this resistance mechanism depends on the support stiffness, with UFC Section 6-9 distinguishing between rigid supports (arching action), non-rigid supports (resistance governed by support stiffness), and simply supported walls (flexural resistance limited by the modulus of rupture of the mortar). The implementation of the damage classification in the parametric P-I sweep is described in Section 4.3.

2.3 Numerical modelling of masonry

The numerical modelling of a masonry wall requires two largely independent choices: how the heterogeneous brick–mortar assembly is discretised, and which constitutive law governs the material response of the brick units. The first choice determines whether the units and joints are represented explicitly or smeared into an equivalent continuum; the second determines how the chosen material develops inelastic deformation and damage. This section surveys the candidate discretisation strategies and the constitutive law relevant to quasi-brittle brick, deferring the selection of a model for this thesis to Section 3.1.

2.3.1 Continuum and discrete representations

Masonry is a composite of stiff units bonded by comparatively weak mortar joints, and its mechanical behaviour is dominated by the joints: cracking, sliding, and debonding concentrate at the unit–mortar interfaces long before the units themselves fail. Numerical strategies for masonry differ chiefly in how faithfully they resolve this interface behaviour. Three families are established in the literature [32]. Macro-modelling smears units and joints into a single homogeneous continuum. Micro-modelling, in its simplified and detailed variants, represents units and joints as distinct entities; a limiting case is the Rigid Body-Spring Model (RBSM), which treats the units as rigid blocks connected by springs at the interfaces. The remainder of this section develops each in turn, followed by the constitutive law for the brick units.

2.3.2 Macro-modelling

In the macro-modelling approach, no geometrical distinction is made between the individual masonry units and the mortar joints. Instead, the wall is treated as a single, homogeneous, continuous material [32]. By mathematically "smearing" the distinct mechanical properties of the constituents into one set of global material parameters, the approach eliminates the need to model the contact interactions at the brick–mortar interfaces.

Macro-modelling is consequently the most computationally efficient strategy for analysing large-scale masonry structures. Its corresponding limitation is an inability to capture localised failure mechanisms: because the material is assumed to remain a continuum, discrete failure modes such as individual mortar crushing, localised brick crushing, and physical brick separation cannot be reproduced. The homogenised description resolves the global deformation of the wall but not the unit–mortar interface failures that govern masonry cracking.

2.3.3 Micro-modelling: simplified and detailed

In contrast to macro-modelling, micro-modelling treats the bricks and the mortar as distinct materials with their respective mechanical properties, explicitly calculating the interactions between them. The strategy is conventionally divided into two sub-categories: simplified and detailed [32].

In a *simplified micro-model*, the bricks are modelled as separate continuum elements capable of local deformation, while the physical thickness of the mortar is not explicitly meshed. Instead, the dimensions of the bricks are artificially expanded to account for the joint, and the mortar is represented as a zero-thickness cohesive interface between the units. These massless interaction layers are assigned cohesive and frictional properties to capture the stress and strain transfer, and the eventual debonding, between the units. The discrete elements are then assembled to act as a global wall. Two examples of this modelling strategy developed within the RISE masonry project are the Abaqus finite element meso-scale model of Godio and Flansbjer [18], in which bricks are governed by a Concrete Damaged Plasticity formulation and zero-thickness cohesive contacts reproduce the mortar joints, and the HybriDFEM model of Bouckaert, Godio, and Almeida [6], which combines finite and discrete elements within a simplified micro-model framework to capture the dynamic out-of-plane response of URM walls. The *detailed micro-model* is a more complex approach. As in the simplified model, each brick is meshed as a separate element; in addition, the mortar joint is explicitly modelled with its true physical thickness using continuum elements possessing mass and volume. The detailed micro-model is the most computationally demanding of the three strategies, but it provides the highest fidelity, capable of resolving a wider array of failure modes within the mortar itself than either the simplified or the macro model.

Between these two variants, the simplified micro-model is widely regarded in the literature as an effective compromise: it captures the failure modes critical to blast analysis, namely mortar joint cracking and individual brick separation, while keeping computational cost manageable. The three strategies are illustrated in Figure 2.1.

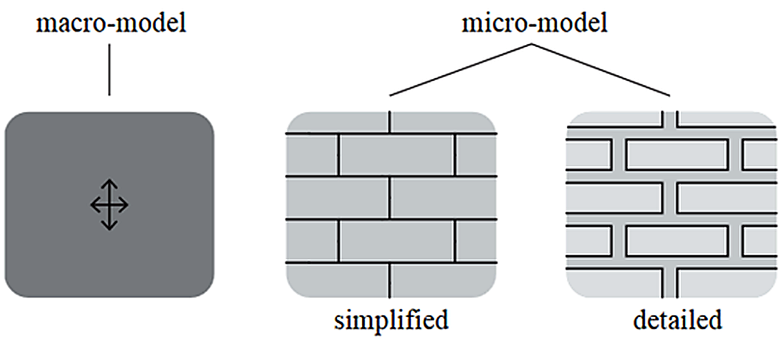

Figure 2.1: Illustration of different masonry modelling strategies: macro-modelling, simplified micro-modelling, and detailed micro-modelling. Reproduced from [21].

2.3.4 Rigid Body-Spring Model

The Rigid Body-Spring Model (RBSM) is a discrete numerical approach originally introduced by Kawai [28] for the limit analysis of structures, including problems involving strong non-linearity and fracture. Rather than discretising the material as a continuum, the RBSM represents the domain as an assemblage of rigid elements connected by springs, with the spring stiffnesses calibrated to recover the elastic response of the underlying material. Crack initiation and propagation are governed by failure criteria assigned to the springs themselves rather than to the rigid bodies, which renders the simulated fracture behaviour objective with respect to the size and geometry of the network [5, 28].

The framework is particularly well suited to masonry. Standard Cauchy continuum approaches struggle to reproduce the macroscopic mode II shear fractures characteristic of in-plane masonry failure [7, 8], whereas the RBSM accounts for the orthotropic shear response of brickwork by allowing distinct shear strengths and friction values to be assigned in directions parallel and normal to the bed joints. This anisotropy enables the model to capture the localised vertical crack patterns and joint-dominated failure modes typically observed in masonry under dynamic loading [7]. Because each unit is reduced to a single rigid node, the global degree-of-freedom count of an RBSM wall is orders of magnitude smaller than that of an equivalent meshed continuum model. The corresponding limitation is that deformation and damage *within* the units themselves, such as tensile cracking and brittle fragmentation of the brick, are not accessible to the model, which confines its validity to loading regimes in which the joints govern the response.

2.3.5 Constitutive modelling of brick units: Concrete Damaged Plasticity

Where the units are modelled as deformable continuum elements, as in both micro-modelling variants, a constitutive law must be assigned to the brick material. The Concrete Damaged Plasticity (CDP) model, sometimes referred to as CDPM, is the constitutive law most widely used for this purpose and represents the current state of the art for quasi-brittle materials. It is a continuum-based model developed to simulate the inelastic behaviour and stiffness degradation characteristic of such materials, combining a plasticity formulation with scalar damage variables that progressively degrade the elastic stiffness.

Although initially formulated for various grades of concrete, including normal concrete and Ultra-High Performance Fibre Reinforced Concrete (UHPFRC) [23], the model has been successfully adapted to the numerical simulation of masonry structures [1, 39]. When implemented in finite-element software such as ABAQUS/Explicit, the model requires the definition of seven material constants alongside tabulated compressive and tensile stress–strain data [39].

In dynamic and extreme-loading scenarios, the CDP model has proven effective for evaluating the response of masonry panels under high-strain loading [39]. Applied to unreinforced and retrofitted brick masonry walls subjected to blast loading, it is capable of capturing the complex inelastic degradation of the masonry [1], and the damage patterns it predicts under blast loading have been shown to correlate closely with experimental observations [1].

2.4 Blast loading

The air through which a blast wave propagates is a highly nonlinear medium. The shock front rises near-instantaneously to a peak overpressure, decays through a positive phase, and is followed by a longer suction phase below ambient [2, 33]. Reflection, diffraction, and clearing of the wave around obstacles are governed by the compressible Navier–Stokes equations and produce loadings on a structure that depend on the geometry of both the source and the surroundings. The fluid-structure interaction then determines how this loading is transferred into the masonry: peak pressures and impulses act on the wall surface, drive its out-of-plane response, and dictate cracking and potential collapse. Approaches to representing this load range from purely empirical curves, through coupled fluid-structure simulation, to fully idealised pulses that omit the air physics entirely [21].

2.4.1 Blast-load representation in the masonry literature

The blast-loaded masonry literature is divided across all three of these families, and no single load idealisation has become standard. Godio et al. [21] review experimental and numerical approaches to the out-of-plane response of URM walls subjected to free far-field blasts, in which the load on the wall is commonly idealised as a spatially uniform far-field overpressure; Parisi, Balestrieri, and Asprone [34] drove a macro-modelled tuff-stone wall with an idealised triangular gas-explosion pulse; and Wei, Huang, and Li [44] assessed brick URM walls against the rotation-based damage criteria of design guidance. The far-field assumption underpinning much of this work is significant: where the standoff is large relative to the wall, the overpressure is approximately uniform across the wall surface and the loading reduces to a single overpressure–time history, the form in which idealised-pulse and empirical blast models are commonly applied [21]. Near-field and contact detonations, in which the fireball and a strongly non-uniform pressure field dominate, fall outside this regime and outside the scope of the load models considered here.

For any P–I representation the pulse must be reducible to a peak pressure and an impulse, and the choice of pulse *shape* is not neutral: the position of the dynamic part of the boundary is sensitive to the pulse shape and rise time [30], and a triangular pulse yields iso-damage curves that differ from those of an exponentially decaying pulse of equal peak and impulse [35]. The three representation families are developed in turn below: the empirical ConWep model, coupled Eulerian–Lagrangian air-domain modelling, and the family of idealised pulses.

2.4.2 Empirical blast loading: ConWep

ConWep (Conventional Weapons Effects) is an empirical model developed by the United States Army Corps of Engineers in the late 1980s, founded on the Kingery–Bulmash equations of 1984 [29]. It produces a predictable pressure-time history that is applied directly to the surface of a structure, with the loading curve obtained by extrapolating from an extensive database of live explosive tests [25]. Two input parameters define the load: the equivalent TNT mass of the charge W and the standoff distance R. The user must additionally specify whether the blast is a free-air burst, in which the wave expands spherically, or a surface burst, in which hemispherical expansion off the ground effectively doubles the energy directed upward.

The model is computationally inexpensive. Because no fluid domain is meshed and no equations of state are integrated, the cost of applying a ConWep load is negligible compared with the cost of the structural analysis itself. Within its designated use cases the loss in accuracy is also small, since the underlying equations are fitted to real explosive test data and the resulting pressure-time history follows the modified Friedlander equation that is generally accepted as the canonical description of a free-field blast wave [2, 15].

2.4.2.1 Hopkinson–Cranz scaling and the Kingery–Bulmash polynomials

ConWep evaluates the blast parameters by exploiting Hopkinson–Cranz scaling, which states that any two blasts producing the same scaled distance generate geometrically similar blast waves [2]. The scaled distance is defined as

$$Z = \frac{R}{W^{1/3}}, \tag{2.4}$$

measured in $\mathrm{m/kg}^{1/3}$, and reduces the two-parameter charge-and-distance space to a single governing variable.

The scaling has consequences for every parameter of the blast wave. Lengths scale proportionally with $W^{1/3}$ at fixed Z, and because blast wave speeds depend only on local pressure (which is invariant at fixed Z), all time quantities scale by the same $W^{1/3}$ factor. Peak overpressure depends only on Z and is invariant under scaling, while specific impulse, the product of pressure and time, inherits the $W^{1/3}$ time scaling [2].

To fully define the blast load at a point, six parameters are required: peak incident overpressure P_{so}, peak reflected overpressure P_r, positive phase duration t_d, incident specific impulse i_s, reflected specific impulse I_r, and shock arrival time t_a [29]. ConWep evaluates each of these internally using high-order polynomials fitted to test data in log-log space:

$$\log_{10} Y = \sum_{n=0}^{N} C_n (K_0 + K_1 \log_{10} Z)^n \,, \tag{2.5}$$

where K_0 and K_1 shift and scale the log-transformed scaled distance into a normalised domain suitable for polynomial evaluation, and C_n are the fitted polynomial coefficients [42]. The number of coefficients N varies between 4 and 14 depending on the parameter; quantities that exhibit a more complex dependence on Z, such as peak reflected overpressure, require higher-order polynomials. Once evaluated, time and impulse parameters are multiplied back by $W^{1/3}$ to recover their dimensional values, while pressure parameters return their dimensional values directly.

2.4.2.2 Load application and the modified Friedlander equation

The Kingery–Bulmash equations supply the peak values and the total duration of the blast pulse but not the pressure at intermediate times. For dynamic structural analysis a complete pressure-time history $P(t)$ is required, and ConWep constructs this using the modified Friedlander equation [2, 15]:

$$P(t) = P_{\text{peak}} \left(1 - \frac{t}{t_d}\right) e^{-bt/t_d}, \tag{2.6}$$

where P_{peak} is the peak pressure (incident or reflected), t_d is the positive phase duration, and b is the dimensionless waveform decay coefficient.

Equation (2.6) describes the full waveform for $t \geq 0$, encompassing both the positive and negative phases. During the positive phase ($0 \leq t \leq t_d$) the linear term $(1 - t/t_d)$ remains positive and ensures that the pressure drops to exactly zero at $t = t_d$, while the exponential term controls the rate at which the pressure decays from its peak: a larger b yields a rapid decay, a smaller b sustains elevated pressures longer. Beyond the positive phase ($t > t_d$) the linear term becomes negative, and when multiplied by the still-positive exponential term it produces a smooth negative phase (suction) that asymptotes to zero from below. In practice the structural loading is dominated by the positive phase, and ConWep implementations typically focus the impulse calibration on this interval.

The decay coefficient b is not provided directly by the Kingery–Bulmash equations and must be determined by requiring that the area under the positive phase of the Friedlander curve equals the corresponding specific impulse from the polynomial fits. The decay coefficient differs between the incident and reflected waves; for the reflected pressure profile the positive-phase area is set equal to the reflected specific impulse I_r:

$$I_r = \int_0^{t_d} P_r \left(1 - \frac{t}{t_d}\right) e^{-bt/t_d} \, dt. \tag{2.7}$$

Integration yields a one-dimensional root-finding problem in b:

$$f(b) = \frac{1}{b} - \frac{1 - e^{-b}}{b^2} - \frac{I_r}{P_r \, t_d} = 0. \tag{2.8}$$

For elements that do not face the source perfectly head-on or edge-on, the angle of incidence θ between the blast wave direction and the element's outward normal must be accounted for. ABAQUS [11] interpolates between full reflected pressure P_r at $\theta = 0°$ and incident pressure P_{so} at $\theta = 90°$ at every time increment using Hyde's empirical cosine-squared formula [25]:

$$P(\theta, t) = P_r(t) \cos^2\theta + P_{so}(t)\Big(1 + \cos^2\theta - 2\cos\theta\Big). \tag{2.9}$$

The $\cos^2\theta$ weighting was derived empirically from experimental pressure measurements at various angles of incidence and can be understood as a consequence of the momentum transfer normal to a surface diminishing with the square of the direction cosine, in a manner analogous to the dynamic pressure component $\frac{1}{2}\rho v^2 \cos^2\theta$ in steady-flow stagnation theory.

2.4.2.3 Limitations

The empirical foundations of ConWep impose a number of limitations. Because no equations of state or Navier–Stokes equations are solved for the air medium, no physical blast wave propagates through the simulation domain; the pressure is mapped and applied directly to the structural surface. As a consequence, reflection, channelling, and shadowing of the wave by surrounding obstacles cannot be represented. ConWep is therefore most accurate for open-field blasts acting on simple, unobstructed geometries.

Near-field and contact detonations are similarly outside the model's scope. ConWep assumes that a fully formed shock front — spherical for a free-air burst, hemispherical for a surface burst — has developed before interacting with the structure. In the near-field the fireball, detonation products, and chemical expansion dominate the pressure regime, and because ConWep contains no physics to generate these phenomena no comparable load profile is produced.

Strictly, the formulation is only valid within the scaled distance range $0.05 \leq Z \leq 40$ [42]. Parameters falling outside this range should be discarded. The lower bound corresponds to the extreme near-field regime adjacent to the fireball, where reliable experimental data does not exist; in this regime the ratio $I_r/(P_r t_d)$ can fall outside the mathematical bounds of the Friedlander integral in Equation (2.8), leaving no physically valid solution for the decay coefficient. The upper bound represents the far-field where overpressures decay to acoustic levels and other formulations become more appropriate.

2.4.3 Coupled Eulerian–Lagrangian air domain modelling

When the limitations of empirical models cannot be overlooked, the alternative is to mesh the air domain explicitly and propagate the blast wave through it using a fluid solver [3]. A purely Lagrangian mesh is unsuitable for this purpose: the nodes move with the material, and under the extreme deformations of a blast wave the resulting elements suffer severe entanglement and hourglassing. The standard remedy is to mesh the air with Eulerian elements, which remain fixed in space while the material flows through them. The pressure-volume-energy relationship of the air and detonation products is then governed by an equation of state, such as the ideal gas law or the Jones–Wilkins–Lee (JWL) equation [31].

In the Coupled Eulerian–Lagrangian (CEL) method the fluid domain (air and explosive charge) is modelled in an Eulerian framework while the structure (the masonry wall) retains its standard Lagrangian formulation [3, 11]. As the Eulerian blast wave propagates and intersects the Lagrangian structural elements, a coupling algorithm — typically a penalty-based contact method — transfers the fluid pressure onto the solid wall. The wall deforms in response, and the new boundary is communicated back to the Eulerian domain so that the air flows around the deformed structure. CEL captures phenomena that empirical models cannot resolve, including localised spalling, breach holes, and the clearing of pressure around fractured masonry blocks.

A related family of methods is the Arbitrary Lagrangian–Eulerian (ALE) formulation, in which the mesh moves independently of both the material and a fixed frame [3]. The mesh is allowed to deform Lagrangianly during a first solver step and is then smoothed and remapped during an Eulerian advection step. ALE is widely used in blast modelling and has comparable capability to CEL for capturing fireball expansion and shockwave propagation.

The cost of CEL relative to ConWep is substantial. The air domain must be meshed at a resolution sufficient to resolve the shock front, which in three dimensions produces element counts an order of magnitude or more above those of the structure alone. Equation-of-state evaluation, contact tracking at the fluid-solid interface, and the explicit time-step constraint imposed by the speed of sound in air all contribute to runtimes that, for a single representative load case, are typically one to two orders of magnitude greater than the corresponding ConWep run.

2.4.4 Idealised pulses

A third approach abandons any physically grounded representation in favour of an idealised pulse: a pressure-time history whose shape is fixed by assumption and whose peak pressure and total impulse are then independent design parameters [30, 35]. In this representation the negative phase is omitted and the positive phase is described by a simple analytic form, applied directly to the loaded surface in the same manner as ConWep. The price paid is that the resulting pulse may not correspond to any physically realisable explosion; the gain is that any (P, I) point in the plane can be specified, and the two parameters are by construction uncoupled.

2.4.4.1 Triangular pulses

The simplest idealised pulse is a triangular pressure-time history with a near-instantaneous rise to peak pressure and a linear decline to zero over a duration t_d. For prescribed peak pressure P_{peak} and total impulse I, the duration is fixed by $t_d = 2I/P_{\text{peak}}$, and the pulse is fully determined. The triangular form has a long history in blast-resistant design and is appealing for its analytical tractability: closed-form SDOF response solutions exist for many resistance functions under triangular loading [2, 30].

Its accuracy as an approximation to a real blast pulse, however, is limited. [35] compared the structural response under triangular and exponentially-decaying pulses of equivalent peak pressure and impulse and showed that the resulting iso-damage curves differ measurably in the dynamic regime of the (P, I) plane. The discrepancy arises because the triangular pulse delivers its impulse more uniformly in time than the exponential pulse, which front-loads the impulse closer to the peak in a manner that more faithfully represents a real Friedlander-shaped blast wave.

2.4.4.2 Exponential pulses

The exponential pulse retains the form of the modified Friedlander equation but treats the parameters as free design variables rather than as quantities derived from a Kingery–Bulmash evaluation. The pulse shape is again

$$P(t) = P_{\text{peak}} \left(1 - \frac{t}{t_d}\right) e^{-bt/t_d}, \tag{2.10}$$

for $0 \leq t \leq t_d$, with the negative phase omitted. The pulse is more faithful to real blast loadings than the triangular form, since it preserves the front-loaded impulse delivery of the Friedlander shape while still allowing P_{peak} and I to be specified independently.

The difficulty is that the exponential pulse has four parameters — P_{peak}, I, t_d, and b — of which only two (P_{peak} and I) are fixed by a prescribed peak pressure and impulse. The remaining two, t_d and b, must be determined together, and there are infinitely many (t_d, b) pairs consistent with any given (P_{peak}, I). A direct application of the Friedlander form therefore requires a supplementary rule for choosing t_d and b, and in the absence of such a rule the load is under-specified.

2.4.4.3 Single-parameter exponential formulation

A reformulation rooted in the Swedish blast-design literature [40, 41] resolves this ambiguity. The modified Friedlander shape is replaced by a pure exponential decay over a semi-infinite positive phase:

$$P(t) = P_{\text{peak}}\, e^{-\kappa t}, \qquad t \geq 0, \tag{2.11}$$

in which κ is a single decay coefficient. Integrating over the positive phase gives a closed-form relationship between κ, the peak pressure, and the total impulse [35, 40]:

$$I = \int_0^\infty P_{\text{peak}}\, e^{-\kappa t}\, dt = \frac{P_{\text{peak}}}{\kappa}, \tag{2.12}$$

so that $\kappa = P_{\text{peak}}/I$ and the pulse is fully specified by the peak pressure and total impulse alone. The positive duration t_d no longer appears as an independent variable: the pulse decays continuously from $t = 0$ rather than terminating abruptly at t_d, and the four-parameter under-determination of the modified Friedlander form is eliminated. With this formulation, P_{peak} and I become genuinely independent inputs: any point in the (P, I) plane can be reproduced by computing the corresponding κ from Equation (2.12), and the load is fully determined.

2.5 Pressure–impulse diagrams

A pressure–impulse (P–I) diagram summarises the response of a structural element under blast as a single curve in the load plane, separating the combinations of peak pressure and impulse that keep the element within a chosen damage state from those that exceed it (the impulse here being the specific impulse per unit area, in kPa·ms). Because the curve corresponds to a specified level of damage rather than necessarily to collapse, such curves are also referred to as damage curves.

2.5.1 Pressure–impulse diagram fundamentals

A pressure–impulse (P–I) diagram is a transformed response spectrum: by re-plotting the maximum response of a system against peak load and impulse rather than against the load-duration-to-period ratio, the impulsive and quasi-static limits appear as a vertical and a horizontal asymptote [30]. Each curve is an iso-damage contour for a fixed structural configuration and a fixed damage level, dividing the plane into a safe region below and to the left and a region of exceedance above and to the right [30]. The curve divides into three regimes — impulsive, dynamic, and quasi-static — corresponding respectively to loads that are over before, comparable to, and long relative to the response time of the structure [10, 30].

2.5.2 Methods for constructing P–I diagrams

The literature on P–I construction can be organised along three largely independent axes: how the curve is *formulated*, how each point is *computed*, and how points are *searched for* [10]. Locating the present method on each axis clarifies what is inherited and what is new.

By *formulation*, P–I curves fall into single analytical expressions, piecewise analytical expressions, and piecewise-mixed formulations in which the asymptotes are given analytically while the dynamic region is obtained numerically and fitted [10]. The approach adopted in this thesis is piecewise-mixed: the pressure and impulse asymptotes are prescribed analytically, while the shape of the dynamic region connecting them is determined by fitting a parametric hyperbola to numerically located damage-threshold points (Section 2.5.7).

By *structural-analysis method*, points are obtained experimentally, from single-degree-of-freedom (SDOF) idealisations, or from continuum finite-element (FE) models [10]. Experimental P–I derivation is rare owing to the cost and hazard of blast testing, so most curves rest on SDOF or FE analysis. Closed-form and energy-balance SDOF solutions are convenient but restricted to simple resistance functions and idealised pulses, and break down once the resistance is strongly nonlinear or several failure modes interact [2, 30]. FE analysis removes these restrictions and captures both global and local failure modes, at the cost of dense three-dimensional meshes and a strong dependence on the calibration of the material model — a dependence that is especially delicate for brittle materials such as concrete and masonry [10]. The present thesis sits at the FE end of this axis, at its micro-modelled extreme.

By *search algorithm*, the dynamic region between the asymptotes is populated in one of two broad ways [10]. The asymptotes are typically obtained analytically from energy balance, after which the dynamic region is either curve-fitted with a closed analytic form — the hyperbolic-tangent-squared relation of Baker et al. [2] or the shifted hyperbola of Oswald and Sherkut, both reported by Krauthammer et al. [30] — or traced directly point by point. Krauthammer et al. [30] review three such schemes: the asymptote-bounded search of Soh and Krauthammer, which evaluates many analyses within shrinking limits around the asymptotes; the threshold search of Ng and Krauthammer, which fixes the pressure and brackets the safe/damaged transition in impulse independently of the asymptotes; and the polar-coordinate bisection of Blasko et al., which casts rays from a pivot point in the failure zone and bisects along each. An earlier predictor–corrector branch-tracing scheme assumes a smooth, continuous curve and can become unstable where that assumption fails [30]. These deterministic searches share recurring weaknesses: they are computationally intensive, generate a considerable amount of data away from the boundary, and are typically tied to a specific element and loading scheme [30]. The more recent literature applies the same unidirectional searches — pressure-controlled, impulse-controlled, and mixed — to FE-derived points, as in the RC-column study of Chernin, Vilnay, and Cotsovos [9].

Within this direct-tracing lineage the schemes differ chiefly in how the next search location is chosen. The asymptote-bounded, threshold, and polar-bisection schemes reviewed above each select it by a fixed rule, a uniform sweep of pressures, impulses, or polar angles, specified in advance of the search. An alternative is to choose the next location adaptively, from the predictive uncertainty of a model fitted to the points found so far, so that evaluations concentrate where the boundary is least well known. Bisection along a fixed or adaptively chosen line locates each individual point in either case; the adaptive variant differs only in how that line is selected. The adaptive sampling this requires is developed in Section 2.5.8.

2.5.3 Influence of multiple failure modes

The fundamentals above describe a single P–I curve, but they carry an implicit restriction: a curve is an iso-damage contour for one failure mode as well as for one damage level [30]. A real structural element can reach its limit state through several distinct mechanisms, and each mechanism has its own characteristic response time and therefore its own curve in the same load plane. A flexural mode, governed by the bending response of the whole element, and a shear mode, governed by the much faster local response at the supports, do not in general share asymptotes, so a load combination that is safe against one can exceed the other [30]. The same reasoning extends to the joint-dominated mechanisms specific to masonry, such as bed-joint sliding and rocking, which respond on different timescales again.

Because the element fails as soon as the first mode is exceeded, the governing boundary is not any one of these curves but their lower envelope: at every point in the (P, I) plane the controlling mode is the one whose curve lies lowest, and the safe region is the intersection of the safe regions of all modes [30]. Which mode controls typically changes along the boundary, so the envelope can switch modes between the impulsive and quasi-static ends and need not be smooth where two mode curves cross. A single finite-element sweep of the kind developed in this thesis traces only the boundary that is critical for the chosen damage metric and wall configuration; it does not, by itself, separate the contributing modes or guarantee that one mode governs across the whole curve. Recovering the full envelope therefore requires either a damage metric sensitive to every mode or a separate sweep per mode.

2.5.4 P–I diagrams for masonry

P–I work is dominated by reinforced-concrete and steel members — beams, columns, slabs, and plates — as the scope of the major reviews and method papers in the field reflects [9, 10, 30]; masonry-specific P–I studies remain comparatively scarce. Wei, Huang, and Li [44] derived P–I diagrams for unreinforced brick masonry walls and supplied the three-tier damage classification adopted in Section 2.2. Parisi, Balestrieri, and Asprone [34] derived P–I diagrams for load-bearing tuff-stone walls failing out-of-plane in a three-hinge flexural mechanism, using a homogenised macro-model in which units and mortar are smeared into an equivalent continuum. That macro-modelling choice cannot resolve the unit–mortar interface failures that govern URM cracking — precisely the mechanism a micro-model is built to capture, and the reason the simplified-micro representation is adopted in this thesis.

2.5.5 Adaptive sampling of expensive limit-state boundaries

Tracing an expensive failure boundary in as few model evaluations as possible is the central problem of surrogate-based active learning in structural reliability and of Bayesian optimisation more broadly [26, 38]. In the reliability setting a surrogate of the limit-state function is built from a small design of experiments and refined by repeatedly evaluating the true model at the single point a learning function identifies as most informative. The AK-MCS method of Echard, Gayton, and Lemaire [13] and the EGRA method of Bichon et al. [4] are the canonical examples: both place the next evaluation where a balance of predictive uncertainty and proximity to the limit state is most favourable, the former through its U-function and the latter through an expected-feasibility function.

The acquisition step of the present sweep (Section 2.5.8) is an instance of the same uncertainty-driven principle.

There is, however, a structural difference between those methods and the present one. AK-MCS and EGRA score a candidate population spread across the entire input space, so their learning functions must explicitly down-weight candidates far from the limit state. The present method instead generates candidates *along* the current boundary estimate, so that weighting is unnecessary and the acquisition reduces to scoring where along the boundary the model is least certain. The lineage is therefore one of shared principle — adaptive, uncertainty-driven evaluation of an expensive boundary — rather than of a transplanted algorithm.

2.5.6 Bisection

Bisection treats the boundary not as a surface to be sampled but as a level set to be located. With either pressure or impulse held fixed, the wall's response is monotone in the other variable, and a single root-finding bisection along that axis converges to a boundary point at logarithmic cost in the bracket width. Two natural variants follow from the choice of which variable to fix: fix-I bisection, in which the impulse is held constant and the failure pressure is found, and fix-P bisection, its dual. Both converge robustly provided the chosen target value lies outside the relevant asymptote: a fix-I bisection at $I \leq c$ cannot succeed because no failure pressure exists, and analogously for fix-P at $P \leq a$.

2.5.7 Parametric curve families

2.5.7.1 Three-parameter hyperbolic model

The simplest functional form consistent with the two-asymptote structure of the boundary is the rectangular hyperbola

$$P(I) = a + \frac{b}{I - c}, \tag{2.13}$$

in which a is the pressure asymptote, c is the impulse asymptote, and b controls the curvature in the dynamic regime. This form has a long history in blast-loading literature, because the single-degree-of-freedom response of an elasto-plastic system under an idealised pulse is well approximated by a boundary of this form [30]. It has three parameters, all with direct physical interpretation.

2.5.7.2 Four-parameter generalised hyperbolic model

The three-parameter form is rigid: once the asymptotes a and c are fixed, the single coefficient b sets both the scale and the sharpness of the transition between them, so a knee sharper or softer than the rectangular hyperbola implies cannot be represented. Because the failure boundary located by the FE sweep is the response of a discrete micro-model rather than of an idealised single-degree-of-freedom oscillator, it is not expected to be a true rectangular hyperbola, and a form with more freedom in the dynamic region is useful purely as an interpolant. The generalised form

$$P(I) = a + \frac{b}{(I - c)^n} \tag{2.14}$$

introduces a fourth parameter n that decouples the curvature of the dynamic transition from the asymptote locations [30]. For $n = 1$ the rectangular hyperbola is recovered, for $n > 1$ the knee is sharper, and for $n < 1$ it is more gradual. The exponent is a fitting degree of freedom, not a physically interpretable quantity, unlike a and c, which retain their meaning as the quasi-static and impulsive capacities, and is introduced solely to let the fitted curve track anchor points that the rectangular form cannot. This flexibility carries a cost: four parameters require at least four boundary points before the fit is identifiable, and n is weakly identified from points confined to the central dynamic region, where its effect on $P(I)$ is small and its contribution away from $I \approx c$ is negligible. Whether the extra parameter is warranted is therefore not assumed in advance but decided from the located points by the model-selection criterion of Section 2.5.7.3.

2.5.7.3 Model selection via the Bayesian Information Criterion

The choice between the two forms can be made on the data themselves rather than fixed in advance. The Bayesian Information Criterion provides a principled basis [37]:

$$\mathrm{BIC} = n_{\text{pts}} \ln\left(\frac{\mathrm{RSS}}{n_{\text{pts}}}\right) + k \ln(n_{\text{pts}}), \tag{2.15}$$

where n_{pts} is the number of boundary points, k is the number of model parameters, and RSS is the residual sum of squares of the fitted model. The criterion penalises additional parameters by $\ln(n_{\text{pts}})$ per parameter, so a more complex model is preferred only if its improved fit overcomes this penalty. The interpretation of $\Delta\mathrm{BIC}$ follows the Kass–Raftery scale [27], on which $\Delta\mathrm{BIC} > 6$ is conventionally regarded as strong evidence in favour of the more complex model.

2.5.8 Adaptive Bayesian sampling

Adaptive sampling decides *where* to place the next boundary point based on the information already collected. Each completed bisection contributes one boundary point; a model is fitted to all such points; and the next bisection target is selected to maximise an acquisition criterion derived from that model. This is the framework of Bayesian optimisation [38] adapted to boundary tracing rather than function minimisation. The fitted model provides a closed-form prediction of the boundary at every I, and reports its own uncertainty, which can be used to direct sampling toward the regions where the curve is least well known.

The posterior over the curve parameters, and hence the predictive uncertainty, is obtained by Markov-chain Monte Carlo. The affine-invariant ensemble sampler of [22], exposed by the `emcee` package [14], is well suited to the moderate dimensionality and mildly non-Gaussian posteriors encountered. Predictive uncertainty is propagated by evaluating the curve at sample-drawn parameters and reporting the empirical standard deviation. The reliability of the posterior summary depends on whether the chain has mixed, diagnosed through the integrated autocorrelation time τ, with the conventional standard that the chain length satisfy $n_{\text{steps}} \geq 50\tau$ for the estimates to be considered well calibrated [14].

3 Methodology

3.1 Masonry model

The masonry model is the first of the three modelling axes that determine the cost and accuracy of a P–I sweep. Unlike the blast load, which serves two distinct purposes and is therefore implemented in two modes (Section 3.2), a single masonry representation is selected and used throughout. The candidate strategies are described in Section 2.3; this section selects among them against the requirements of the thesis. The selection is made in two passes along two independent axes: first the discretisation strategy, namely how the brick–mortar assembly is divided into units and joints; then the representation of the brick unit itself, namely whether it is a rigid body or a deformable continuum body.

3.1.1 Requirements for the masonry model

Two requirements, pulling in opposite directions, govern the choice of masonry model. The first is fidelity to the failure physics. The intended output of the tool is a P–I failure boundary, and the damage metric on which that boundary is based is the global out-of-plane support rotation of the wall (Section 4.3). For a blast-loaded unreinforced masonry (URM) wall, this global response is itself mediated by the mortar joints: cracking, sliding, and debonding at the unit–mortar interfaces are the mechanisms through which the wall deforms and rotates. The model must therefore represent the mortar joints as discrete interfaces capable of cracking and sliding, since it is their degradation that produces the global flexural response the sweep measures. A representation that cannot resolve the joints as distinct entities cannot reproduce the physics behind the support rotation.

The second requirement is computational economy. A single P–I sweep entails on the order of tens to hundreds of finite-element evaluations (Section 3.3), and the model must remain cheap enough for the sweep to complete within a practical wall-clock budget. A representation that is faithful but prohibitively expensive per evaluation is as unsuitable for a sweep tool as one that is cheap but cannot resolve the governing joint behaviour. The selected model must satisfy both requirements simultaneously.

3.1.2 Rejection of macro-modelling

Macro-modelling (Section 2.3.2) satisfies the second requirement comfortably: smearing the units and joints into a single continuum makes it the cheapest of the candidate strategies. It fails the first. Because the homogenised continuum carries no explicit unit–mortar interface, the mortar joints have no discrete representation at all, and the cracking and sliding through which a URM wall mobilises its out-of-plane resistance cannot be reproduced. The smeared description resolves the global deformation of an equivalent homogeneous panel, but not the joint-mediated mechanism that governs the support rotation of a real masonry wall. For a tool whose output depends on that mechanism, the loss of fidelity is decisive, and macro-modelling is not adopted.

3.1.3 Simplified versus detailed micro-modelling

With the continuum approach rejected, the discretisation strategy is drawn from the micro-modelling family (Section 2.3.3), which by construction represents the units and joints as distinct entities and therefore satisfies the fidelity requirement. The choice is then between its simplified and detailed variants.

The detailed micro-model meshes the mortar with its true physical thickness using continuum elements possessing mass and volume. It offers the highest fidelity, but the additional mesh and the contact resolution it demands make it the most expensive of the strategies per evaluation, a cost that is multiplied across every point of a P–I sweep and is difficult to justify when the joint behaviour of interest can be captured without meshing the mortar explicitly.

The simplified micro-model represents the mortar as a zero-thickness cohesive interface between units whose dimensions are expanded to absorb the joint. It resolves the joint cracking and sliding that govern the global flexural response, without the cost of a meshed mortar layer. It is also the representation most commonly adopted in the contemporary blast-loaded masonry literature, and selecting it places the present work within established research practice rather than on an idiosyncratic footing. The simplified micro-model is therefore adopted as the discretisation strategy for this thesis.

3.1.4 Unit representation: rigid versus deformable

A second choice, independent of the discretisation strategy, concerns how the brick unit itself is represented. Within a micro-model the unit may be treated either as a rigid body, as in the Rigid Body-Spring Model (Section 2.3.4), or as a deformable continuum body. The Rigid Body-Spring Model was considered first and was attractive on computational grounds, since representing each brick by a single rigid node carrying six degrees of freedom, rather than by a meshed deformable body, reduces the global degree-of-freedom count by orders of magnitude.

An RBSM-based wall was not, however, carried through to the final framework. The attempt to construct one in ABAQUS/Explicit met implementation limitations that made the approach impractical within the solver; these are documented in Section 4.1. The unit was therefore modelled as a deformable continuum body instead, and the constitutive law for that body is selected in Section 3.1.5.

3.1.5 Constitutive model for the deformable unit

With the unit modelled as a deformable continuum body, a constitutive law must be assigned to it. Concrete Damaged Plasticity (CDP), described in Section 2.3.5, is adopted for this purpose. The choice follows established research practice: CDP is the state-of-the-art constitutive law for quasi-brittle materials and has been applied repeatedly to unreinforced and retrofitted masonry under blast loading, with predicted damage patterns shown to correlate closely with experiment [1, 39]. To capture the strain-rate sensitivity of masonry strength under blast loading, the CDP definition is supplemented with a Dynamic Increase Factor (DIF); the implementation and verification of the rate-dependent hardening are described in Section 4.1.

3.1.6 Composed selection

The masonry model adopted for the tool is a simplified micro-model. The mortar joints are represented as zero-thickness, surface-based cohesive interfaces between the units, and the brick units are modelled as deformable continuum bodies governed by the Concrete Damaged Plasticity constitutive law with rate-dependent hardening. Macro-modelling is excluded because its homogenised continuum carries no discrete joint interface and so cannot reproduce the joint-mediated mechanism that governs the global support rotation on which the sweep relies; the detailed micro-model is excluded because meshing the mortar explicitly is too expensive for a tool that must evaluate the model hundreds of times in a sweep. The brick unit is modelled as a deformable continuum body rather than a rigid one, the Rigid Body-Spring Model having proven impractical to implement in ABAQUS/Explicit. The resulting model satisfies both the fidelity and the economy requirements of Section 3.1.1, and is the representation most commonly used in the contemporary blast-loaded masonry literature. The implementation of this model in ABAQUS, including the trials that led to it and the solver-level difficulties encountered along the way, is described in Chapter 4.

3.2 Blast load modelling

The blast load is the second of the three modelling axes that determine the cost and accuracy of a P-I sweep. Unlike the masonry model, where a single representation is selected and used throughout, the blast load model serves two distinct purposes in this thesis, and the GUI exposes both as user-selectable modes. In *validation mode* the tool reproduces a specific real explosion event with a known charge weight and standoff distance, allowing the wall response to be compared against an instrumented or published experiment. In *sweep mode* the tool places probes at arbitrary (P, I) coordinates chosen by the boundary-tracing algorithm of Section 3.3, irrespective of whether the resulting pulse corresponds to any physically realisable explosion.

These two modes pull in opposite directions, and no single load model serves both well. Each mode therefore has its own load representation implemented behind the GUI. The candidate models — the empirical ConWep formulation, coupled Eulerian–Lagrangian (CEL) air-domain modelling, and the family of idealised pulses — are described in Section 2.4; this section selects among them for each mode, beginning with the physically grounded models for validation and proceeding to the idealised pulses for the sweep itself.

3.2.1 Selection for validation mode

The first methodological question addressed in this thesis was which of the two physically-grounded load models, ConWep (Section 2.4.2) or CEL (Section 2.4.3), should back the GUI's validation mode. Validation mode reproduces real, instrumented or published experiments in which a charge of known mass is detonated at a known standoff from the wall, and the comparison is made against a measured response. Both models can in principle represent such a scenario; the question is which does so more efficiently for the geometric scope of this thesis. Of the two Eulerian-type approaches discussed in Section 2.4.3, only CEL was tested against ConWep in this work; ALE is referenced there for completeness but was not exercised.

The masonry model exposed by the tool is a single wall with no surrounding terrain, adjacent structures, roofs, or internal compartmentalisation. In this geometry the capabilities that distinguish CEL from ConWep — the explicit modelling of reflection, channelling, shadowing, and near-field fireball physics — have no surface on which to act. None of the phenomena CEL can resolve and ConWep cannot are present in the validation geometry implemented in the GUI. The cost differential, by contrast, is decisive: a ConWep run is essentially free relative to the structural analysis, while CEL multiplies the total runtime by one to two orders of magnitude with no corresponding gain in accuracy for this class of problem.

ConWep was therefore selected as the load model behind validation mode, and CEL is not exposed in the current version of the tool. Should the scope of the masonry model be extended in future work to include adjacent walls, terrain features, roofs, or other obstacles, the reflection-capturing capability of CEL would become relevant and adding it as a third selectable load model would be a natural extension. This is discussed further in Chapter 8.

3.2.2 Suitability of ConWep for sweep mode

The selection above settles the validation question. Sweep mode, however, is a different problem. The sweep requires not a faithful representation of any one explosion but a load whose peak pressure and total impulse can be specified independently, so that the algorithm of Section 3.3 can place probes at chosen (P, I) coordinates. This requirement is not obviously incompatible with ConWep: by varying the charge mass W and standoff distance R, a continuous family of (P, I) pairs can be generated, and the question is whether this family covers enough of the (P, I) plane to support a usable failure boundary.

To answer this, ConWep was used to map the (P, I) pairs reachable across the full valid scaled-distance range $0.05 \leq Z \leq 40$ [42] for a representative wall geometry. The resulting domain is examined in Section 6.4 in the discussion, where it is compared against the (P, I) region in which the failure boundary of interest lies.

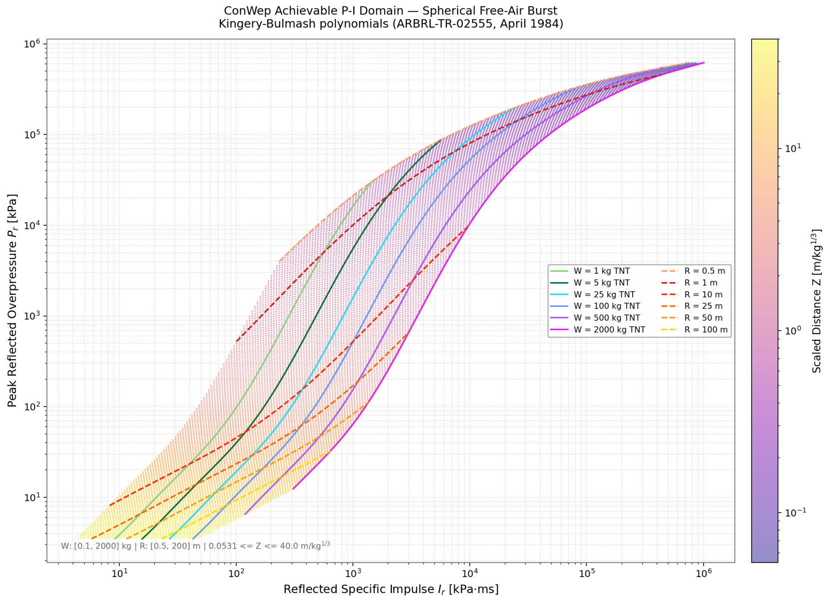

Figure 3.1: Achievable P-I domain through ConWep.

3.2.2.1 Coverage and the quasi-static regime

The conclusion of that mapping is that ConWep does not cover enough of the (P, I) plane for a P-I sweep to be feasible. The reachable domain is bounded above by the near-field cutoff at $Z = 0.05$ and tapers off in the far-field; more critically, it is largely absent from the quasi-static region of long impulse and low pressure. This absence is not a defect of ConWep so much as a feature of the physics it represents: real explosions producing very long-duration positive phases at low overpressure are uncommon, since such loadings would require either very large charges at very large standoffs or scenarios that are not free-air bursts at all (confined detonations, deflagrations, or sustained pressure events). The Kingery–Bulmash database [29] does not contain such cases and the empirical fits do not extrapolate into them.

For a failure boundary that is intended to span from impulsive through dynamic to quasi-static loading, the absence of coverage in the quasi-static regime is disqualifying. The asymptotic structure of the boundary, in particular the location of the pressure asymptote a, cannot be resolved if no probes can be placed at long impulse. ConWep is therefore unsuitable as the load model behind sweep mode.

3.2.2.2 Coupling between P and I

A second, more fundamental issue arises from the way ConWep parameterises the load. The peak pressure and total impulse of a ConWep pulse are both functions of the same two inputs, W and R, and cannot be varied independently. Choosing a (P, I) target therefore amounts to inverting two coupled functions, and the inverse mapping is constrained to lie on the manifold of physically realisable (W, R) combinations. Even in regions where the domain is reachable, the algorithm cannot place a probe at an arbitrary (P, I) coordinate; it can only ask for the nearest realisable pair. This coupling alone would make ConWep awkward as a sweep load even if its coverage were complete.

3.2.3 Selection of the sweep load model

With ConWep ruled out, the sweep load is drawn from the idealised pulses of Section 2.4.4, which satisfy by construction the two requirements ConWep fails: full coverage of the (P, I) plane and independent control of P and I. Of the three forms, the triangular pulse departs measurably from the more realistic Friedlander-shaped response and yields iso-damage curves that differ in the dynamic regime [35], which is undesirable for a sweep whose objective is an accurate failure boundary rather than a conservative design envelope. The two-parameter exponential pulse is more faithful but is under-determined for a (P, I) target, requiring a supplementary rule to fix t_d and b.

The single-parameter exponential formulation of Section 2.4.4.3 resolves both difficulties. Setting $\kappa = P_{\text{peak}}/I$ through Equation (2.12) specifies the pulse completely from the two sweep targets alone, while retaining the front-loaded impulse delivery of the Friedlander shape. It is therefore the load model implemented behind the GUI's sweep mode: it covers the entire (P, I) plane without restriction to a Kingery–Bulmash domain, and its P and I are uncoupled by construction.

3.2.4 Composed selection

The GUI exposes two operational modes, each backed by the load model selected for its purpose. In validation mode, ConWep [25, 29] is applied directly to the loaded surface using the charge weight and standoff distance supplied by the user; the open-field, single-wall geometry of the masonry model lies entirely within ConWep's regime of empirical validity, so none of the phenomena that distinguish CEL from ConWep can manifest, and the cost differential against CEL is not justified for this class of problem. In sweep mode, the single-parameter exponential pulse [40, 41] is applied with the user-supplied (P, I) target driving the decay coefficient through Equation (2.12); ConWep is unsuitable for this mode because its (P, I) coverage is incomplete in the quasi-static regime and its inputs are coupled, while the triangular pulse departs measurably from the more realistic exponential shape [35] and the two-parameter exponential pulse is under-determined. The two modes share the same wall model and damage criteria; they differ only in how the loading on the wall surface is constructed. The mechanics of the mode toggle, including how the GUI translates user inputs into the corresponding ABAQUS load definition, are described in Chapter 4. Extensions of the masonry model that would justify exposing CEL as a third load option are discussed in Chapter 8.

3.3 P–I sweep algorithm

Once the masonry and blast-load models of the preceding chapters have been fixed, every evaluation of a candidate (P, I) point reduces to a single ABAQUS run, and the cost of constructing the diagram is set entirely by how those evaluations are placed. Two properties of the problem dictate the design. First, every evaluation is expensive: a single ABAQUS Explicit run for a representative wall geometry takes minutes to hours depending on mesh density, so any scheme that wastes evaluations far from the boundary is impractical. Second, the response near the asymptotes is degenerate: at $I \leq c$ no finite pressure produces failure and at $P \leq a$ no finite impulse does, so an algorithm that does not respect the asymptotes will repeatedly probe regions where the wall cannot fail by construction. The algorithm must therefore trace the boundary as tightly as possible in as few evaluations as it can while remaining robust to these degeneracies. It is built from the four components introduced in Section 2.5: bisection, a warm-start scheme, a hyperbolic fit with promotion, and Bayesian uncertainty to drive probe placement.

3.3.1 Bisection as the locating primitive

A dense tensor-product grid would resolve the boundary, but at a cost that grows quadratically with resolution and is spent overwhelmingly on points far from the curve; bisection (Section 2.5.6) instead locates a boundary point directly along a single axis, the basis of the direct-tracing "advanced search" algorithms in the P-I literature [10]. Both fix-I and fix-P variants are used — the impulse- and pressure-controlled searches of prior deterministic schemes [9, 10] — the direction being chosen per iteration as described below. Ray bisection, in which a ray is cast from the origin, was not adopted: on the log-log plot natural to a P-I diagram a ray from the origin is no longer a straight line, and its parametrisation must be carefully normalised if the bisection is not to favour one axis disproportionately, which forfeits the conceptual simplicity that motivates it.

Because a fix-I bisection at $I \leq c$ and a fix-P bisection at $P \leq a$ cannot converge, the algorithm tracks the current asymptote estimates at all times and selects only target values that lie outside the relevant asymptote. The asymptote estimates are taken from the current parametric fit, which couples the bisection direction and target to the rest of the algorithm rather than fixing them in advance.

3.3.2 Warm start

The sweep begins with a strategic seed phase: a small number of fix-I and fix-P bisections are run over wide brackets to produce an initial set of boundary points with no prior model to guide them. These cold bisections are comparatively expensive in evaluations, but they are few, and they bootstrap the first parametric fit.

Every subsequent bisection is warm-started. Two sources of prior information narrow the bracket before the root-find begins. The first is the current fitted curve: its closed-form prediction of the boundary pressure at the chosen impulse gives a tight initial estimate around which the bracket can be centred. The second is proximity to already-found boundary points: when a new probe lies near an existing anchor, that anchor's pressure bounds the bracket from one side, reducing the upper or lower limit and removing the failed/survived region that the existing point has already resolved. The combined effect is that a warm-started bisection begins from a narrow interval and converges to a new anchor in a small number of evaluations, in contrast to the wide-bracket seed bisections.

3.3.3 Hyperbolic fit and promotion

The boundary points are described by the parametric hyperbolae of Section 2.5.7. Each sweep begins with the three-parameter rectangular hyperbola of Equation (2.13), which is identifiable from few points and whose three parameters carry direct physical meaning. Promotion to the four-parameter generalised form of Equation (2.14) is not assumed in advance but triggered from the data, using the Bayesian Information Criterion of Equation (2.15): the four-parameter model is adopted only once $\Delta\mathrm{BIC} > 6$ on the Kass–Raftery scale, the threshold conventionally read as strong evidence for the more complex model.

This conservative promotion rule reflects the cost of the extra parameter noted in Section 2.5.7. The exponent n is identifiable only once enough points populate the dynamic region, and fitting the four-parameter model prematurely inflates the posterior uncertainty without improving the fit. Beginning with the three-parameter form and promoting only on a clear BIC margin therefore avoids spurious curvature on sparse early data while still allowing the sharper or softer dynamic transitions that the rectangular hyperbola cannot represent.

3.3.4 Probe placement by predictive uncertainty

The parameters of the active hyperbola, and the predictive uncertainty of the fitted curve, are estimated by Markov-chain Monte Carlo using the affine-invariant ensemble sampler exposed by `emcee` (Section 2.5.8). The chain length and walker count are chosen so that the mixing criterion $n_{\text{steps}} \geq 50\tau$ is met in typical operation, where the observed integrated autocorrelation time for the four-parameter model reaches $\tau \approx 60$.

Candidate impulse values for the next probe are not drawn from the whole (P, I) plane but generated *along the current fitted boundary*, sampled at uniform arc length. This restriction to the boundary is the key structural difference from full-space active-learning reliability methods such as AK-MCS and EGRA [4, 13], in which a learning function (the U-function or expected feasibility) must explicitly down-weight candidates far from the limit state because the candidate pool spans the entire input space. Here that weighting is unnecessary: every candidate already lies on the boundary estimate, so the acquisition need only score *where along the boundary* the model is least certain. The present method shares the uncertainty-driven principle of those methods while operating on a boundary-restricted candidate set rather than a Monte Carlo population.

Each candidate is scored by the predictive standard deviation of the fitted curve in $\log_{10} P$ space,

$$\alpha_{\text{var}}(I) = \sigma_{\log_{10} P}(I). \tag{3.1}$$

The logarithmic measure is used in place of the linear predictive standard deviation because the linear posterior variance of a hyperbola diverges near the impulse asymptote — small perturbations to the asymptote location c produce large changes in P — so a linear max-variance criterion would be dominated by the asymptote irrespective of where information is most useful. The relative $\log_{10} P$ standard deviation is bounded and comparable across the multi-decade pressure range.

The uncertainty score alone, however, has two weaknesses on this problem. First, it concentrates probes near the impulse asymptote where $\sigma_{\log_{10} P}$ peaks, starving the quasi-static region. Second, and more fundamentally, the parametric hyperbola is not expected to be a faithful model of the true failure line: the boundary of a micro-modelled URM wall need not follow a rectangular or generalised hyperbola, so the fit carries an irreducible structural error and its predictive variance reflects only how tightly the *assumed* form has been pinned, not how well that form matches the wall. Acquisition driven purely by this variance would therefore over-trust a possibly misspecified model. Both weaknesses are mitigated by blending the uncertainty score with a coverage score $\alpha_{\text{cov}}(I)$, the log-space distance from a candidate to its nearest existing boundary point, through a coverage-bias weight $w \in [0, 1]$:

$$\alpha(I) = (1 - w)\,\hat{\alpha}_{\text{var}}(I) + w\,\hat{\alpha}_{\text{cov}}(I), \tag{3.2}$$

where $\hat{\alpha}_{\text{var}}$ and $\hat{\alpha}_{\text{cov}}$ are each normalised to $[0, 1]$. The default $w = 0.7$ weights coverage more heavily than uncertainty, deliberately treating the fitted curve as a sampling scaffold rather than a trusted predictor: it yields an approximately even arc-length distribution of anchors along the boundary, which is the desired outcome when the curve itself cannot be relied upon. Lowering w recovers more uncertainty-driven exploration for cases where the parametric form is known to be adequate. Once at least three points are present, an additional arc-length *zone-balancing* step partitions the boundary into zones and directs the next probe into the most under-populated zone before maximising $\alpha(I)$ within it, further guarding against the systematic crowding near the asymptote that a global $\arg\max$ over the model variance would otherwise produce. The selected impulse then defines the location for the next warm-started fix-I bisection of Section 3.3.2, closing the loop between fitting and sampling.

3.3.5 Composed method overview

The four components compose into a single algorithm. The sweep opens with the seed phase, in which a small number of cold fix-I and fix-P bisections produce an initial set of boundary points. A three-parameter hyperbolic model is fitted to these points by MCMC. The fit is then refined adaptively: at each iteration the blended acquisition function of Equation (3.2), subject to arc-length zone-balancing, selects an impulse value; a warm-started fix-I bisection establishes a new boundary point there; and the model is refitted, promoting to the four-parameter form once the BIC criterion is met. The loop terminates once a user-specified number of boundary points has been located, subject to an evaluation-budget ceiling as a safeguard.

Termination is governed by point count: the sweep stops once the user-specified number of anchor points has been located, subject to the evaluation-budget ceiling as a safeguard. A stopping rule based on the fit's own convergence was not adopted, because the fitted hyperbola is a scaffold rather than a faithful model of the boundary, so its convergence reflects only how tightly the assumed form has been pinned and not whether the boundary is well traced. Fixing the number of anchors instead places the resolution of the boundary under the user's direct control and makes the cost of a sweep predictable, with the fitted curve serving only as an interpolant between the located points.

No single component is novel: bisection is classical, the BIC is half a century old, MCMC is a textbook tool, and adaptive Bayesian sampling has been studied extensively in the optimisation [26, 38] and structural-reliability [4, 13] literatures. The methodological contribution lies in the composition.

3.4 Modelling trials and implementation

The models selected in the preceding sections were not arrived at directly. Several approaches were trialled and abandoned during development, and the final model is the product of that process. This section records the trials in chronological order. The principled selections are stated in Sections 3.1, 3.2, and 3.3; the purpose here is to document the implementation route and the ABAQUS-specific obstacles that shaped it.

3.4.1 Two-brick model and the cohesive interface workaround

The modelling process began with a two-brick model in *ABAQUS*, inspired by previous work conducted at RISE that simulated the well-known van der Pluijm experiments [16]. The primary objective of this initial dynamic model was to evaluate the behaviour of the mortar joint, specifically its non-linear response to tension and shear, modeled as an interaction between rigid bricks. Initially, parameters such as tensile strength, compressive strength, and fracture energy were set to arbitrary values with the intention of calibrating them later.

Although *ABAQUS* restricts natively modeling instances as rigid bodies, this was bypassed by initially creating deformable blocks and then applying rigid body constraints to lock their inertia, effectively forcing them to act as rigid blocks. To represent the mortar, a surface-based interaction was initially applied between the bricks.

However, a significant difficulty was encountered in transferring the interface model into the explicit dynamic solver. The frictional and cohesive interaction definitions functioned without issue in the *ABAQUS/Standard* (implicit) solver, but when the same surface-based cohesive definition was carried over to *ABAQUS/Explicit*, the software issued a warning and forced the interaction into a purely linear response. This bypassed the intended cohesive damage and failure mechanics of the mortar entirely, which was critical for capturing the joint degradation observed in the van der Pluijm tests.

To circumvent this restriction, a workaround was required to effectively 'trick' the explicit solver. Instead of relying solely on surface-based interaction properties, an additional physical layer consisting of cohesive elements was introduced between the bricks. By modeling the mortar as this separate, discrete cohesive layer, the explicit solver was able to correctly process the cohesive degradation and properly simulate the non-linear failure of the interface under dynamic conditions. Once the explicit cohesive interface was established, it was necessary to verify that the bricks would physically separate as intended under extreme dynamic conditions. To isolate the interface behaviour and verify the failure criteria without the compounding variables of a full blast load, a simplified kinematic test was conducted.

Rather than applying a complex pressure wave, the structural elements were subjected to a high initial velocity in the out-of-plane direction. This controlled test allowed for the direct observation of the cohesive layer's damage evolution and verified whether the rigid blocks would properly detach once the mortar's fracture energy was fully dissipated.

The cohesive element layer described here was a workaround specific to this early two-brick trial. It was later superseded: as development progressed, the mortar was modelled directly as a surface-based cohesive interface assigned to a General Contact definition, with cohesive stiffness, a maximum-stress damage initiation criterion, and energy-based damage evolution specified as a contact interaction property rather than through a discrete element layer. This surface-based cohesive interface is the mortar representation carried into the final framework, and it is the form described in Chapter 4 and used in the case study.

3.4.2 Scaling limitations and the shift to deformable meso-modelling

While the initial two-brick Rigid Body Spring Model (RBSM) successfully simulated interface separation, scaling this approach to a full wall revealed significant limitations within *ABAQUS*. Defining hundreds of individual brick instances as discrete rigid bodies proved exceedingly tedious, and rigid body constraints often caused the entire wall assembly to act unnaturally as a single fused unit rather than independent blocks.

To bypass this restriction, an intermediate approach was tested where the bricks were modeled as deformable elements utilizing the Concrete Damage Plasticity (CDP) model, but assigned an artificially high elastic modulus to mimic rigid behaviour. However, it was quickly deduced that this workaround defeated the primary computational purpose of the RBSM approach. If the explicit solver was already expending computational resources to calculate stress states within deformable CDP elements, there was no longer a benefit to artificially stiffening them. Consequently, the purely rigid body approach was abandoned in favor of a detailed meso-scale approach, assigning the bricks their proper, realistic material parameters. This trial is the practical basis for the unit-representation choice made in Section 3.1.4.

3.4.3 Boundary conditions and pre-stress strategy

With the transition to a fully deformable wall model, the next step was to apply realistic boundary conditions and pre-stressing, namely gravity and axial loads, prior to the blast application. Two strategies for establishing the pre-stressed state were trialled before the one carried into the final framework was settled on.

The first was a sequential implicit-to-explicit approach. The intention was to simulate the static pre-stressing in an implicit *ABAQUS/Standard* model and then use the `IMPORT` function to transfer the established stress states into the explicit dynamic blast model. Despite extensive troubleshooting, this proved unsuccessful for the masonry assembly. While global stresses were imported, the contact stresses and cohesive states at the discrete mortar interfaces failed to transfer correctly between the solvers. At this stage of development the mortar was still represented through the intermediate cohesive element layer described above, and the cohesive states carried by that layer did not survive the implicit-to-explicit transfer. The sequential import strategy was therefore abandoned.

The second strategy was to keep the analysis within a single solver and use the native restart capability of *ABAQUS/Explicit*: a quasi-static equilibrium analysis would establish the pre-stressed state, and the blast analysis would then be restarted from the converged end of that step. Technically this transferred correctly, the contact and cohesive states surviving the restart because no solver change is involved. At the time, however, the blast load was still represented through ConWep, whose blast reference points must be defined in the equilibrium step. Fixing the reference points in the equilibrium step makes the standoff distance a property of that step, which cannot then be varied per evaluation, so a restart-based ConWep sweep could not place probes at different standoffs. The restart approach was therefore set aside while ConWep remained the load model.

The restart strategy was reinstated once the sweep load model was changed to the single-parameter exponential pulse (Section 3.2). The exponential pulse is applied as a surface pressure and requires no blast reference point, which removes the constraint that had made restart incompatible with the sweep. With that obstacle gone, the restart architecture became usable: a single equilibrium model is solved once and every blast evaluation is restarted from its converged state. This is the approach adopted in the final framework, described in Section 4.1.1.

3.4.4 Element formulation diagnostic

The transition from validation to the parametric sweep called for a model of lower computational cost, since the sweep evaluates the wall across many load cases rather than a single scenario. Reduced-integration C3D8R solid elements were therefore adopted for the sweep, being the less expensive of the available continuum hexahedral formulations.

The first sweep runs exhibited severe and unphysical mesh distortion under high-impulse loading, causing analyses to abort prematurely. Three contributing causes were identified through separate diagnostics. The first was an overly aggressive mass-scaling factor: the scaling applied to keep the explicit time increment economical had been set high enough to inflate the element inertia artificially, corrupting the dynamic response. The second was the boundary condition configuration: a fixed–fixed out-of-plane support produced very large local deformations near the constrained edges, whereas a pinned out-of-plane support reduced these substantially by relieving the edge restraint. These two causes were identified by trialling fully integrated C3D8 elements as a diagnostic. As shown in Figure 3.2(a), the C3D8 mesh distorted severely under a representative high-impulse case ($P = 50$ kPa, $I = 3000$ kPa · ms), demonstrating that the distortion was not specific to the reduced-integration formulation and pointing instead to the mass scaling and boundary conditions as underlying causes.

The third cause was the absence of enhanced hourglass control on the C3D8R elements. Reduced-integration elements are susceptible to zero-energy deformation modes, which hourglass control suppresses. A separate diagnostic was conducted using the original C3D8R formulation with enhanced hourglass control activated, all other settings unchanged. As shown in Figure 3.2(b), this alone was sufficient to produce a physically stable response under the same high-impulse loading, with no spurious deformation.

The final framework therefore adopts C3D8R elements with enhanced hourglass control, a corrected mass-scaling factor, and a pinned out-of-plane support. All three corrections are necessary and are carried into every sweep evaluation.

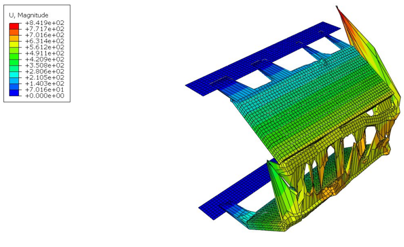

(a) Severe numerical distortion with fully integrated C3D8 elements.

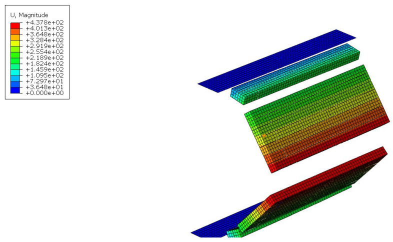

(b) Stable response with reduced-integration C3D8R elements and enhanced hourglass control activated.

Figure 3.2: Element formulation diagnostic under high-impulse blast loading ($P = 50\ \mathrm{kPa}$, $I = 3000\ \mathrm{kPa \cdot ms}$). The fully integrated C3D8 mesh distorts severely and aborts (a), confirming that the distortion was not isolated to the reduced-integration formulation. Activating enhanced hourglass control on the C3D8R formulation, with all other settings unchanged, produces a physically stable response (b). Together with the corrected mass-scaling factor and pinned boundary conditions identified separately, this combination is carried into the final framework.

4 Framework

This chapter presents the framework that integrates the modelling components into the final automated tool. Whereas the preceding chapters selected each model in isolation, this chapter describes the dynamic masonry model, the blast load, and the P-I sweep as they function in the final product, together with the graphical user interface that combines them into a single workflow. It sets out the adopted dynamic model and the verification of its rate-dependent behaviour, the form in which the blast load is applied, the damage results extracted from each analysis, the implementation of the P-I sweep, and the GUI through which the complete pipeline is operated.

4.1 Dynamic model

The dynamic model adopted in the final tool is a deformable simplified micro-model. Each brick unit is meshed as a deformable continuum body governed by the Concrete Damaged Plasticity constitutive law with rate-dependent hardening, discretised with C3D8R reduced-integration solid elements with enhanced hourglass control. The mortar joints are represented as zero-thickness surface-based cohesive interfaces between the units.

The static pre-stress, namely gravity and the axial pre-compression, is established in a explicit quasi-static equilibrium model, and the subsequent dynamic blast load is applied in a separate explicit step restarted from the converged equilibrium state. This restart architecture is described in Section 4.1.1; the development route that led to it, including the approaches trialled and abandoned along the way, is documented in Section 3.4.

4.1.1 Equilibrium model and restart architecture

The wall is not rebuilt for each sweep evaluation. Instead, a single equilibrium model is constructed once and solved to a converged pre-stressed state, and every blast evaluation is then launched as a restart from that state. The equilibrium model establishes the wall geometry, the material assignments, the cohesive mortar interfaces, the boundary conditions and the static pre-stress, namely gravity and the axial pre-compression. It is run until the kinetic energy has decayed to a negligible fraction of the internal energy, the check described in Section 4.5.5, so that the state passed to the blast analysis is genuinely in static equilibrium. This converged state is the common origin shared by every evaluation in the sweep.

Each sweep evaluation restarts the analysis from the end of the equilibrium step and applies the blast pulse of Section 4.2 as the only quantity that differs between evaluations. Because the equilibrium model, which is the expensive part to construct, is solved only once and reused, the per-evaluation cost is reduced to that of the blast step alone, and the pre-stressed state entering every blast analysis is guaranteed identical. The route by which this restart architecture was arrived at, and the alternative pre-stress strategies trialled before it, are documented in Section 3.4.

4.1.2 Implementation and verification of strain-rate effects

Following the establishment of the full-scale explicit model, it was critical to accurately capture the material's dynamic hardening under blast loads. In explicit dynamic events, masonry strength is highly dependent on the strain rate. The Dynamic Increase Factor (DIF) was implemented into the CDP material definition. The DIF for the brick's ultimate compressive strength was calculated using the bilinear regression of Hao and Tarasov [24], as reported by Rafsanjani, Lourenço, and Peixinho [36]:

$$\mathrm{DIF} = \begin{cases} 0.0268 \ln \dot{\varepsilon} + 1.3504 & \text{for } \dot{\varepsilon} \leq 3.2\ \mathrm{s}^{-1} \\ 0.2405 \ln \dot{\varepsilon} + 1.1041 & \text{for } \dot{\varepsilon} > 3.2\ \mathrm{s}^{-1} \end{cases} \tag{4.1}$$

No experimental data is available in the literature for the strain-rate dependence of the tensile strength of clay brick masonry; this limitation was also encountered by Rafsanjani, Lourenço, and Peixinho [36], who applied the same DIF for tension, shear, and compression in the absence of dedicated experimental data. The same simplifying assumption is adopted in the present work: the tensile DIF is taken equal to the compressive DIF given by Equation 4.1. Theoretical DIF values for strain rates ranging from 1 to $1000\ \mathrm{s}^{-1}$ were calculated and tabulated for both tension and compression.

To confirm that *ABAQUS* correctly interpreted and applied these tabulated strain-rate dependencies, an isolated single-brick verification model was developed. The brick was meshed with C3D8R elements and assigned the CDP material definition. The base of the brick was encastred, and a prescribed compressive velocity of $500\ \mathrm{mm/s}$ was applied to the top face, producing a state of approximately uniform uniaxial compression. The analysis was run in *ABAQUS*/Explicit with a step time of $0.01\ \mathrm{s}$. Two simulations were run under identical boundary conditions and identical prescribed velocity, differing only in the assigned material: one with the tabulated rate-dependent hardening active (DIF active), and one with the equivalent rate-independent hardening (DIF disabled).

The stress–time and strain-rate–time histories extracted from both simulations are presented in Figure 4.1. In both cases the prescribed velocity produced an equivalent plastic strain rate that fluctuated about approximately $\dot{\varepsilon} \approx 10\ \mathrm{s}^{-1}$, with instantaneous values ranging between roughly 7 and $16\ \mathrm{s}^{-1}$ (Figures 4.1b and 4.1d). The two strain-rate histories are closely comparable, confirming that any difference in the stress response is attributable to the material definition alone and not to a difference in loading. The contour of equivalent plastic strain (PEEQ) at the end of the analysis, shown in Figure 4.2, is uniform across the brick in both cases, confirming that the specimen deforms homogeneously and that the recorded element-level stresses represent a genuine constitutive response rather than a localised numerical artefact.

The Mises stress responses differed substantially in magnitude. The DIF-disabled simulation rose to a peak corresponding to the static compressive ultimate strength of approximately $13.4\ \mathrm{MPa}$ before post-peak softening (Figure 4.1c). The DIF-active simulation, subjected to the identical loading, reached a peak of approximately $22\ \mathrm{MPa}$ before softening (Figure 4.1a). Both responses exhibit the same characteristic rise-to-peak followed by post-peak softening, differing only in magnitude.

This magnitude difference is consistent with the expected rate-scaled response. At the observed strain-rate range, Equation 4.1 gives a dynamic increase factor of approximately 1.66–1.70 (for example, $0.2405\ln(10) + 1.1041 \approx 1.66$). Applied to the static compressive ultimate strength of $13.4\ \mathrm{MPa}$, this predicts a dynamic ultimate strength of approximately $22\ \mathrm{MPa}$, in close agreement with the simulated DIF-active peak. Equivalently, the ratio of the two simulated peak stresses, $22.0/13.4 \approx 1.64$, agrees with the theoretical dynamic increase factor to within a few percent. This isolated test verified the DIF implementation — both its correct construction in the material table and its correct application by *ABAQUS* — prior to applying the air-blast pressure wave to the full wall. It should be noted that this is a verification of the implementation only; the DIF values themselves are taken from the literature regression of Equation 4.1 and are not independently validated against experimental data for the masonry considered here.

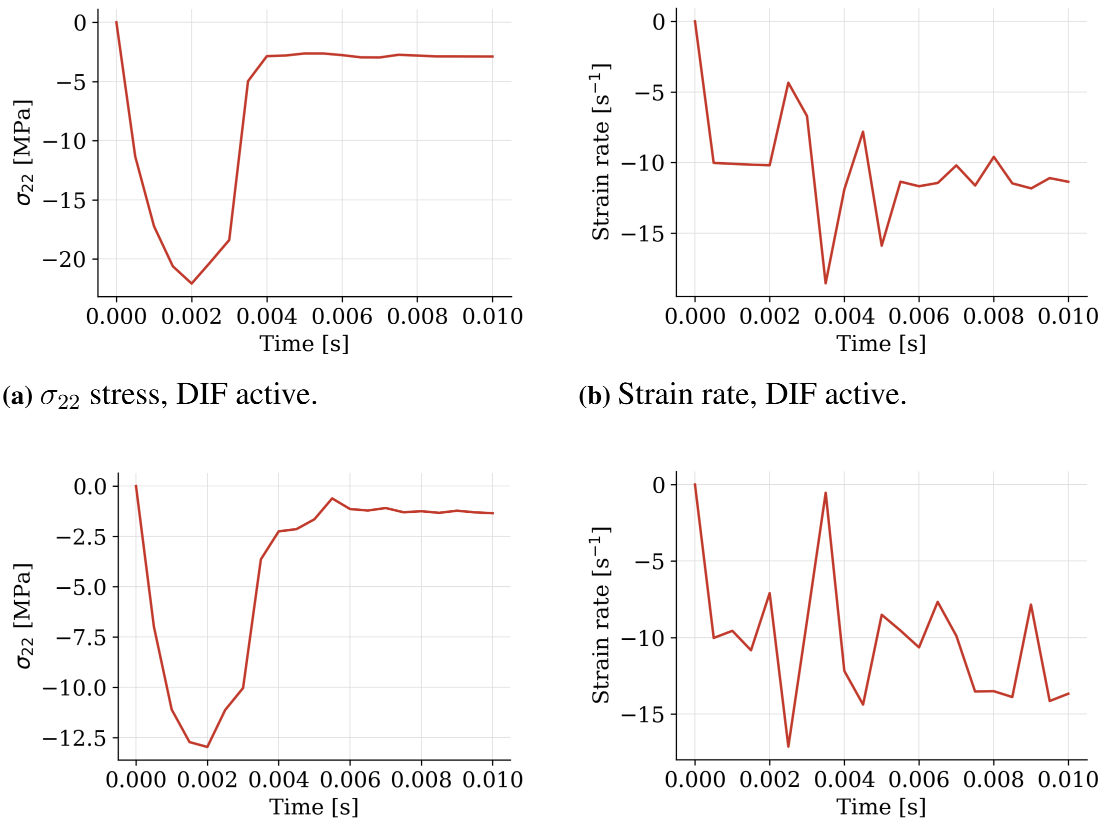

(a) σ_{22} stress, DIF active.

(b) Strain rate, DIF active.

(c) σ_{22} stress, DIF disabled.

(d) Strain rate, DIF disabled.

Figure 4.1: Single-brick DIF verification under prescribed compressive velocity ($500\ \mathrm{mm/s}$ at the top face, base encastred). The strain-rate histories (b, d) fluctuate about $\dot{\varepsilon} \approx 10\ \mathrm{s}^{-1}$ in both cases, while the axial stress response σ_{22} (a, c) differs in magnitude between the DIF-active and DIF-disabled simulations. The DIF-active peak ($\approx 22\ \mathrm{MPa}$) and the DIF-disabled peak ($\approx 13.4\ \mathrm{MPa}$, the static compressive ultimate strength) give a ratio of ≈ 1.64, in close agreement with the theoretical DIF of ≈ 1.66 at the observed strain rate. Stress in MPa, strain rate in s^{-1}, time in seconds.

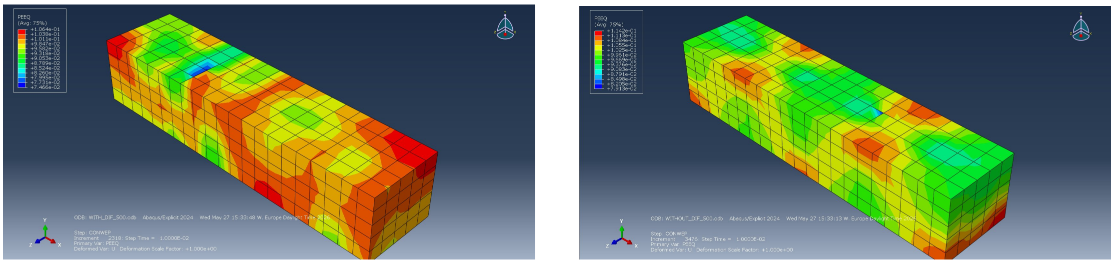

(a) PEEQ, DIF active.

(b) PEEQ, DIF disabled.

Figure 4.2: Contour of equivalent plastic strain (PEEQ) at the end of the single-brick analysis. The plastic strain field is uniform across the brick in both the DIF-active (a) and DIF-disabled (b) cases, confirming homogeneous deformation and that the recorded element-level stresses represent a genuine constitutive response rather than a localised numerical artefact.

4.2 Blast application

The blast load applied in the final tool is the single-parameter exponential pressure pulse selected in the methodology. In sweep mode the pulse $P(t) = P_{\text{peak}}\, e^{-\kappa t}$, with decay coefficient $\kappa = P_{\text{peak}}/I$ set from the target peak pressure and impulse, is applied as a uniform pressure over the front face of the wall. In validation mode the load is instead applied through ConWep from a user-specified charge weight and standoff distance.

Within the framework, the blast load is the single quantity that varies between sweep evaluations. The wall geometry, the material model, the boundary conditions, and the pre-stress are all held fixed by the equilibrium model, and each sweep evaluation differs from the next only in the (P, I) pair driving the pulse. This separation is what allows the dynamic model and the blast load to be combined under the control of the sweep algorithm without rebuilding the wall for every load case.

4.3 Damage assessment

4.3.1 Global damage: support rotation extraction

Global damage is assessed through the out-of-plane chord rotation θ, defined in Equation 2.3. The wall height H is determined programmatically by scanning the undeformed nodal coordinates of all brick instances in the model and recording the maximum y-coordinate. Midspan nodes are identified geometrically at $y = H/2$.

Time histories of out-of-plane displacement (U_1) are extracted at the support nodes and midspan nodes from the Abaqus output database (`.odb`) for both the peak dynamic phase (`Blast_Step`) and the residual post-blast phase (`Post_Blast_Step`). The peak midspan deflection δ_{mid} and the corresponding chord rotation θ_{mid} are then compared against the three-tier damage classification of Wei, Huang, and Li [44] described in Section 2.2. The residual phase is retained to capture cases where the wall survives the peak loading but remains in a non-reusable state.

4.3.1.1 Conservative treatment of the non-reusable bound

The framework of Wei, Huang, and Li [44] formally defines collapse only through the deflection criterion $\delta \geq t$, where t is the wall thickness. The rotation thresholds $\theta_{\text{fail1}} = 1.0°$ and $\theta_{\text{fail2}} = 0.5°$ delineate the boundaries between reusable, non-reusable, and the remaining performance space without explicitly classifying walls that exceed the non-reusable rotation bound. For the present work, a conservative classification convention is adopted: walls whose peak rotation exceeds the non-reusable threshold of $1.0°$ are treated as collapsed, even though the deflection-based collapse criterion is not evaluated. The rationale is that a wall whose rotation has surpassed the highest serviceability limit defined in the literature can no longer be considered fit for its intended structural function. Classification is therefore based on the peak midspan rotation alone, which is the damage metric extracted reliably across the sweep.

The resulting classification rule, implemented in the post-processing stage of the P-I sweep, is:

- **Reusable**: $\theta_{\text{mid}} < 0.5°$.
- **Non-reusable**: $0.5° \leq \theta_{\text{mid}} < 1.0°$.
- **Collapse**: $\theta_{\text{mid}} \geq 1.0°$, OR the simulation aborted due to excessive element distortion.

The second collapse condition, simulation abort on excessive distortion, captures cases where the explicit solver fails to complete the analysis due to severe local deformation, which is itself a physical indicator that the wall has lost coherent structural response. Such cases are classified as collapse without further post-processing, since the rotation history becomes unreliable beyond the point of distortion-induced abort.

4.3.2 Output and visualisation

The damage metrics are written to a structured CSV file immediately following each analysis, and the corresponding `.odb` file is deleted to preserve disk space during the pressure-impulse (P-I) sweep. Each record holds the peak and residual values of θ and δ and the resulting damage-state classification, which together provide the boundary points from which the P-I diagram is constructed.

4.4 Sweep implementation

The P-I sweep repeatedly evaluates the dynamic model of Section 4.1 under the exponential blast pulse of Section 4.2, classifying each evaluation by the damage criterion of Section 4.3. Each evaluation is a single ABAQUS/Explicit analysis discretised with the C3D8R reduced-integration elements; the diagnostic that established this element choice over the fully integrated C3D8 formulation is reported in Section 3.4.4.

The sweep is driven by the adaptive Bayesian algorithm described in the methodology, organised as a two-level architecture. An orchestrating driver holds the state of the sweep, namely the boundary points located so far and the current parametric fit, and applies the acquisition logic to select the next pressure–impulse evaluation point. Each selected point is dispatched to an independent worker analysis, which builds and runs the corresponding blast case and returns its damage metrics to the driver. Separating the two levels keeps the expensive finite-element evaluation isolated from the lightweight decision logic: the driver never holds an analysis in memory, and a failed or aborted evaluation is contained to its own worker without interrupting the sweep.

After each evaluation the damage metrics are recorded and the analysis output database is discarded to limit disk usage, as described in Section 4.3. The accumulated results form the running record from which the driver refits the boundary and selects the following point. Because that record is written incrementally, the sweep is checkpointed by construction: an interrupted run can resume from the last completed evaluation rather than restarting from the beginning. The orchestration mechanics and the resumption behaviour are described as part of the workflow summary in Section 4.5.5.

4.5 Graphical user interface

A central objective of this thesis was to develop a graphics-assisted tool that enables researchers and engineers to perform automated non-linear FEM modelling of blast-loaded masonry structures without requiring direct modification of source code. The tool is implemented as a parametric graphical interface integrated with Abaqus/CAE, which generates the wall geometry, material assignments, boundary conditions, and analysis configuration from a series of structured input dialogs, and then drives the adaptive P-I sweep that those inputs configure.

4.5.1 Design requirements

The following requirements guided the design of the tool:

1. **Parametric wall generation.** The user should be able to specify brick dimensions, bond pattern, number of courses, and boundary conditions without editing source code. The tool should derive all dependent quantities (wall height, length, and thickness) automatically.

2. **Persistent profiles.** Wall configurations should be storable as named profiles that can be recalled, compared, and reused across analysis sessions.

3. **Sweep configuration.** The user should be able to configure the Bayesian adaptive P-I sweep parameters — pressure and impulse bounds, evaluation budget, and target boundary points — from the same interface.

4. **Compatibility.** The tool must function within the constraints of managed university workstations, where administrative privileges are unavailable and plug-in installation directories may be restricted.

4.5.2 Parametric input

The interface collects all parameters through a sequence of dialog pages, grouped so that no single page presents an overwhelming number of fields. The pages cover the brick dimensions; the wall layout and mesh density; the bond pattern; the boundary conditions; the Bayesian sweep settings; and the output preferences such as the profile name and project directory. The mortar joint is not exposed as an input: the simplified micro-model represents it as a zero-thickness surface-based cohesive interface rather than as a meshed layer, so the joint has no thickness for the user to set.

Once a configuration is entered it is stored as a named profile. Profiles persist between sessions and can be recalled or compared, so that a wall defined once need not be re-entered, and one profile is designated as the active configuration used to build the model.

4.5.3 Bond pattern options

Three bond patterns are available, applied when the wall assembly is generated:

- **Running bond** (stretcher bond): each course consists of full-length bricks with a half-brick stagger between adjacent courses. This is the most common bond pattern in unreinforced masonry construction and produces the strongest interlocking between courses.

- **Stack bond**: all courses are identical with no stagger, so the vertical joints align across courses. This produces a geometrically regular but structurally weaker pattern, and is included to allow comparative study of the effect of joint alignment on blast resistance.

- **Flemish-style face pattern**: full-length and half-length bricks alternate along each course, with adjacent courses offset so that a half brick is centred on the full brick below it. A structural Flemish bond proper requires a wall at least two wythes thick, since it relies on headers spanning the wall thickness; the single-wythe model used here reproduces only the alternating face appearance and the associated joint-stagger geometry, not the through-thickness header interlock.

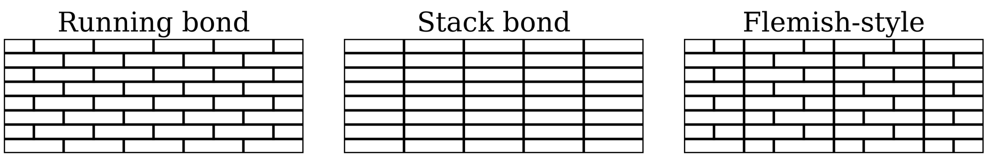

Figure 4.3: The three bond patterns available in the tool: running bond (left), stack bond (centre), and Flemish-style face pattern (right).

The wall length is computed as

$$L_z = n_{\text{full}} \cdot l_{\text{full}} + l_{\text{edge}} + n_{\text{full}} \cdot g_{\text{mortar}} \tag{4.2}$$

where n_{full} is the number of full bricks per course, l_{full} and l_{edge} are the full and edge brick lengths, and g_{mortar} is the mortar gap (zero in this implementation). The wall height is

$$H_y = n_{\text{layers}} \cdot (h_{\text{brick}} + g_{\text{mortar}}). \tag{4.3}$$

4.5.4 Boundary condition options

Five boundary condition configurations are available, each defined by the constraints applied to rigid reference points coupled to the top and bottom wall surfaces. The options are summarised in Table 4.1. The selected configuration is carried through the analysis automatically, including the transition from the equilibrium step to the blast step, in which the precompression is released and the vertical displacement of the top of the wall is locked.

Table 4.1: Boundary condition options available in the GUI. Coordinate system: u_1 is out-of-plane (blast direction), u_2 is vertical, u_3 is in-plane horizontal. All BCs are applied to rigid reference points coupled to the top and bottom wall surfaces.

Option	**Bottom RP**	**Top RP**
Fixed–Fixed	Fully fixed	u_1, u_3 locked; u_2 free; all rotations locked
Simply Supported	u_1, u_3 locked; u_2 free; all rotations free	u_1, u_3 locked; u_2 free; all rotations free
Fixed–Pinned	Fully fixed	u_1, u_3 locked; u_2 free; all rotations free
Cantilever	Fully fixed	u_3 locked; u_1, u_2 free; all rotations free
Pinned about Z	u_1, u_2, u_3 locked; ur_1, ur_2 locked; ur_3 free	u_1, u_3 locked; u_2 free; ur_1, ur_2 locked; ur_3 free

4.5.5 Workflow summary

The complete analysis workflow using the graphics-assisted tool proceeds in five stages:

1. **Parametric input.** The user opens the input interface from within Abaqus/CAE and works through the dialog pages, which collect the wall geometry, bond pattern, boundary conditions, sweep parameters, and output preferences. The configuration is saved as a named profile.

2. **Model generation.** The tool reads the active profile and constructs the simplified micro-model assembly, assigning materials, contact properties, and boundary conditions, and submits the quasi-static equilibrium analysis.

3. **Equilibrium verification.** The user confirms that the equilibrium analysis has settled, by checking that the kinetic energy has decayed to a negligible fraction of the internal energy by the end of the equilibrium step.

4. **Adaptive sweep.** The user launches the P-I sweep, which adaptively selects pressure–impulse evaluation points and runs one blast analysis per point. Results are checkpointed after every evaluation, so that an interrupted sweep can be resumed without repeating completed work.

5. **Post-processing.** The boundary points and fitted curve produced by the sweep are assembled into the final P-I diagram.

This workflow reduces the barrier to entry for parametric blast analysis of masonry walls: the user interacts only with dialog boxes, without modifying source code or constructing Abaqus model files directly.

5 Case study

To demonstrate the framework end-to-end and to assess whether the developed tool produces a physically reasonable P–I diagram, a single case study was performed on a representative one-way spanning unreinforced masonry (URM) wall. The case study follows the workflow described in Chapter 4: a wall profile is defined through the GUI, the equilibrium FE model is built and solved, and the adaptive Bayesian P–I sweep is launched, with the driver orchestrating one worker analysis per evaluation. The results of every run are collected in a structured log, and the reusable–non-reusable boundary is fitted and exported by the same driver.

5.1 Isolating a single failure mode

A P–I diagram of a real masonry wall is not a single curve but a family of curves, one per failure mode (out-of-plane flexure, in-plane shear, joint sliding, sliding shear at the supports, and so on), as discussed in Section 2.5.3. In a complete protective assessment, each mode produces its own iso-damage boundary in the P–I plane, and the governing boundary at a given $(P,\ I)$ point is the envelope of these curves. A single sweep of the FE model traces whichever boundary is critical for the chosen damage metric and the chosen wall configuration, and cannot, by itself, distinguish between modes.

The case study presented here is therefore deliberately scoped to a single failure mode, *out-of-plane flexure*, and the wall profile is designed to make that mode dominant. Three design choices contribute to this isolation:

- **Slenderness.** The wall is chosen with a height-to-thickness ratio of $H/t = 24$, placing it in the slender regime in which the out-of-plane response of a URM panel is expected to be governed by flexural deformation rather than by shear at the supports or by compressive crushing [12, 21].

- **Out-of-plane boundary conditions.** The top and bottom edges are modelled as simply supported (pinned) in the out-of-plane direction. The supports therefore carry no out-of-plane moment, and the in-plane translational degree of freedom is left unrestrained at the supports, so that no in-plane membrane stiffness can develop in the model and the resistance is mobilised in the FE model entirely through flexural deformation.

- **In-plane boundary conditions.** The same edges are fixed in the two in-plane directions. Rigid-body sliding of the wall as a whole and in-plane shear deformation at the supports are therefore kinematically prevented in the FE model.

Together, these choices push the flexural P–I boundary downwards (the wall is made comparatively weak in bending) while pushing the boundaries of the other modes upwards (the wall is made comparatively strong in shear and sliding), with the intent of leaving flexure as the first mode to be triggered as the load is increased. This separation is not guaranteed in an absolute sense: a sufficiently impulsive load can excite higher-order modes that the chosen configuration does not entirely suppress, and a systematic check for the activation of competing failure modes was not performed in the present work. The P–I curve reported in this chapter should therefore be read as the *flexural* reusable boundary for the case-study wall under the chosen modelling assumptions, not as the governing boundary for an arbitrary load case.

5.2 Wall configuration and modelling parameters

The case-study wall represents a single-wythe URM panel of the geometry implied by the slenderness and boundary conditions described above. The full set of parameters used for the case study is listed in Table 5.1.

Table 5.1: Case-study wall profile.

Parameter	Value
Wall height H	1200 mm (24 courses $\times$ 50 mm)
Wall length L	1155 mm (5 $\times$ 210 mm + 105 mm)
Wall thickness t	50 mm (single wythe)
Slenderness H/t	24
Brick unit (full length, width, height)	$210 \times 50 \times 50$ mm
Mortar joint thickness	0 mm (zero-thickness cohesive contact)
Bond pattern	Running
Boundary condition (out-of-plane)	Pinned–pinned
Boundary condition (in-plane)	Fixed–fixed

Material properties are those described in Section 4.1: the bricks are modelled with the Concrete Damaged Plasticity (CDP) constitutive law calibrated with rate-dependent hardening, and the mortar joints are represented as zero-thickness surface-based cohesive contacts with the Van de Pluijm experimental parameters ($t_{n,0} = 0.62$ MPa, $G_f = 0.058$ N/mm). Bricks are meshed with C3D8R reduced-integration solid elements with enhanced hourglass control to suppress the volumetric locking and spurious distortion observed in the C3D8 trials reported in Section 3.4.

The blast load is applied as the exponential pressure pulse $P(t) = P_{peak}e^{-\kappa t}$ with $\kappa = P_{peak}/I$, applied as a uniform surface pressure on the front face of the wall. Gravity and the static pre-compression are introduced first in a quasi-static equilibrium model. The blast load is then applied in a separate explicit step restarted from that converged equilibrium state, as described in Section 4.1.

5.3 Sweep configuration

The adaptive sweep was configured to locate the boundary between the *Reusable* and *Non-reusable* damage classes, corresponding to a midspan support rotation $\theta_{\text{mid}} = 0.5°$ following UFC 3-340-02 (Table 6-2). The pressure and impulse search ranges spanned $P \in [2,\, 50]$ kPa and $I \in [10,\, 500]$ kPa·ms, chosen to cover the full range of low-rise blast loading typical of the protective design problems for which the tool is intended. The driver was allocated a maximum evaluation budget of 100 jobs, initialised with 4 seed points distributed in the parameter space, and instructed to locate 9 anchor points on the reusable boundary. The Bayesian acquisition function selected each new candidate point so as to reduce the residual uncertainty in the inferred boundary, and the sweep terminated once either the requested 9 anchor points had been located or the evaluation budget was exhausted. The wall-clock time-out per Abaqus evaluation was set to 6 hours ($21{,}600$ s).

The sweep terminated after 60 completed evaluations, having located the requested 9 anchor points to within the prescribed uncertainty tolerance; the remaining 40 jobs of the budget were not consumed. The distribution of damage classifications across the probed points is summarised in Table 5.2.

Table 5.2: Distribution of damage classifications across the 60 sweep evaluations.

Damage class	**Count**	**Peak θ_{mid} range**
Reusable	27	$0.06° - 0.50°$
Non-reusable	25	$0.50° - 0.98°$
Collapse	8	$4.59° - 43.12°$

The sweep comprising these 60 finite-element evaluations completed in approximately 35 hours of wall-clock time.

5.4 P–I diagram

The algorithm produces two outputs from the case-study sweep: a log of every evaluation, with its applied pressure and impulse, status, peak midspan rotation, and assigned damage class; and a hyperbolic fit of the reusable boundary. The final fit takes the four-parameter form

$$P(I) = a + \frac{b}{(I - c)^n} \tag{5.1}$$

with parameters $a,\ b, c, n$. The fitted boundary, together with the 60 probed load combinations, is shown in Figure 5.1. The nine anchor points the driver selected as straddling the reusable–non- reusable transition are listed in Table 5.3.

Table 5.3: Anchor points selected by the Bayesian driver as straddling the reusable–non-reusable boundary.

Impulse $[kPa \cdot ms]$	**Pressure** $[kPa]$
55.43	39.06
57.31	10.54
60.12	9.38
64.81	7.81
96.68	4.66
161.33	3.23
255.04	2.98
400.00	3.01
500.00	2.89

The fitted boundary returns a quasi-static pressure asymptote of $a \approx 1.33$ kPa, an impulsive impulse asymptote of $c \approx 55.38$ kPa·ms, and a dynamic transition between the two asymptotes spanning approximately $60 \lesssim I \lesssim 250$ kPa·ms.

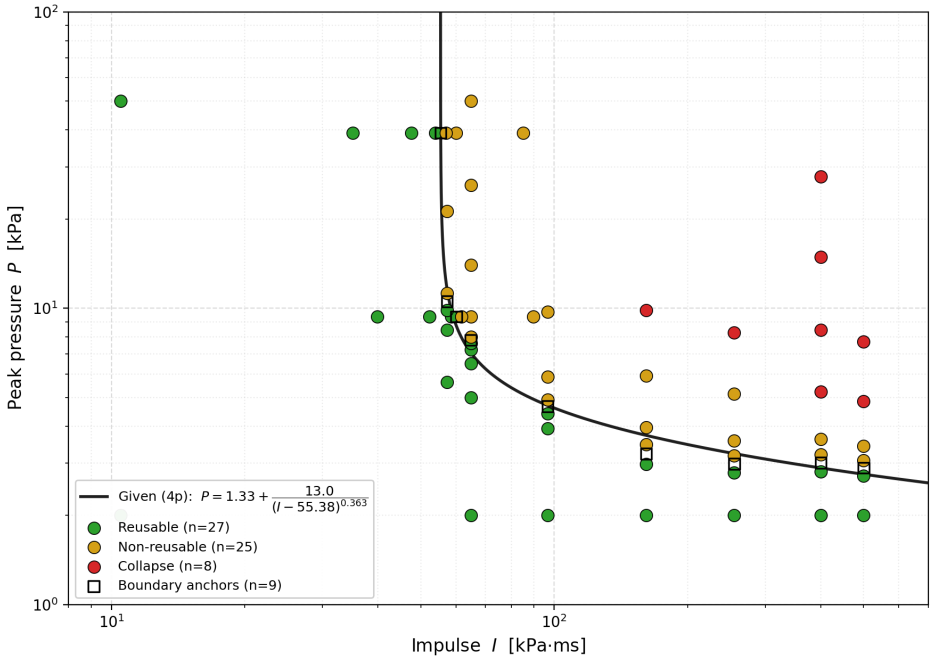

Figure 5.1: Reusable boundary fit produced by the algorithm. The four-parameter hyperbola is overlaid on the 60 probed load combinations, coloured by damage class (Reusable: $\theta_{\text{mid}} < 0.5°$; Non-reusable: $0.5° \leq \theta_{\text{mid}} < 1.0°$; Collapse: $\theta_{\text{mid}} \geq 1.0°$ or excessive element distortion).

5.5 Damage classification across the parameter space

The damage distribution shown in Table 5.2 is informative beyond the boundary itself. The reusable and non-reusable classes together account for 87% of all evaluations, while the eight collapse events occur exclusively at high impulse ($I \geq 161$ kPa·ms) and at pressures well above the reusable boundary. The latter were probed as part of the algorithm's initial seed-point exploration of the upper region of the P–I plane; the mechanisms by which the subsequent evaluations were drawn towards the boundary region are discussed in Section 3.3.4.

A clear gap is visible in θ_{mid} between the non-reusable and collapse classes: the non-reusable runs all peak in the narrow range $0.50°$–$0.98°$, while every collapse event records $\theta_{\text{mid}} \geq 4.59°$, with five of the eight cases exceeding $40°$. No values intermediate between these bands are observed in the case-study sweep.

6 Discussion

The case study demonstrates that the coupled components developed in this thesis operate together to produce a closed P–I failure boundary for an unreinforced masonry wall. Sixty finite-element evaluations sufficed to locate nine anchor points straddling the reusable–non-reusable transition, with the remaining forty jobs of the allocated budget left unconsumed. In that purely operational sense, the framework satisfies the design goal set out in Section 1.3. The single four-parameter fit the program reports is not, however, the only boundary consistent with the data. Two alternative analyses of the same anchor points performed in this chapter (a three-parameter regression and a piecewise interpolation) yield curves that differ from the reported fit in non-trivial ways, with the pressure asymptote in particular bracketed only between roughly 1.33 and 3.05 kPa depending on the form chosen. Beyond the choice of fit, the resolution of the boundary in pressure is bounded from below by the bisection tolerance imposed at each anchor, the coverage of the sweep is bounded by the parameter ranges and budget configured by the user, and the absolute values of the asymptotes are bounded above by the absence of external validation. The remainder of this chapter discusses these features in turn.

6.1 Physical plausibility of the boundary

The boundary recovered from the sweep exhibits the qualitative shape predicted by the simplified single-degree-of-freedom theory: a finite pressure asymptote at long impulse, a finite impulse asymptote at high pressure, and a smooth hyperbolic transition between them. No kinks or discontinuities indicative of a competing failure mode crossing the dynamic regime are visible, and the damage classifications in Table 5.2 cluster tightly around the boundary rather than scattering across the parameter space. To the extent that a P–I diagram for a one-way spanning URM wall under flexural loading is expected to look like a smooth two-asymptote hyperbola, the curve produced here meets that expectation.

What the curve cannot do is validate itself. The asymptotes returned by the sweep, a pressure asymptote near 3 kPa and an impulse asymptote near 55 to 60 kPa·ms, cannot be compared against a published benchmark, because the case-study wall is a fabricated configuration rather than a documented experimental specimen. Both values are nonetheless of a plausible order: a pressure asymptote of a few kilopascals is consistent with the lateral pressure needed to overcome the gravity-induced arching resistance of a slender URM panel, and an impulse asymptote of tens of kPa·ms with the momentum threshold of the same wall under near-instantaneous loading. Beyond such order-of-magnitude checks the boundary remains unverified.

The strongest assurance available within the present framework is that each anchor point is a directly located boundary point: the bisection that produces it brackets the reusable to non-reusable transition along its axis to within the per-bisection tolerance, independently of any fitted curve. This makes the located points themselves reliable to that tolerance, but it does not validate the boundary they trace. The parametric fit, as argued in Section 3.3.4, is only an interpolant between these points and carries no independent claim to correctness, and the points themselves are confined to a single fabricated wall with no measured response to check them against. The external validation required to close this gap, repeating the sweep against an instrumented blast test of the same wall in validation mode, lies outside the scope of this thesis and is identified as future work in Chapter 8.

6.2 Interpretation of the reported boundary

The program reports a single failure boundary as its final output: the four-parameter algorithmic fit of Figure 5.1. Promotion from the three-parameter form to the four-parameter form was triggered automatically by the BIC criterion of Section 2.5.7.3 once the sweep had collected sufficient anchor points to make the additional curvature parameter statistically warranted. Within the logic of the algorithm, this is the boundary the tool was designed to deliver.

In one specific sense the fit succeeds. The additional curvature parameter n permits the dynamic transition through the knee of the curve to be matched smoothly through the dense cluster of anchor points between 60 and 100 kPa·ms, with no discontinuities and no large residuals against the anchor data. The reported curve passes near every anchor and reproduces the qualitative two-asymptote shape predicted by simplified SDOF theory. As a smooth surrogate over the dynamic regime, the four-parameter fit is internally consistent.

The fit is less successful in two other respects. First, its pressure asymptote is an extrapolation rather than a measured quantity. The fitted value, $P_0 \approx 1.33$ kPa, lies below every reusable anchor probed in the sweep, the lowest of which sits at $P \approx 2.89$ kPa (Table 5.3). With no probes below it, the asymptote is fixed primarily by the curvature parameter $n \approx 0.36$, not bracketed by data. Its tight posterior should not be read as confidence in its location: it measures how tightly the four-parameter form is constrained by the points, not how well the true physical asymptote is known. Second, the fit imposes a single smooth global shape on a boundary that may contain local features, such as transitions between sub-modes or contact-state irregularities, that no global function can reproduce. The BIC promotion of Section 3.3.3 certifies the four-parameter form as statistically warranted by the data; it does not certify that the form is physically correct.

To examine whether these limitations are intrinsic to the data or artefacts of the four-parameter form, two further analyses are performed in this section on the same anchor points.

Three-parameter regression as a check on the asymptotes.

The first analysis steps back to the three-parameter rectangular hyperbola $P(I) = P_0 + \dfrac{b}{I - I_0}$, the form the sweep began with before the BIC criterion promoted it to the four-parameter model (Section 3.3.3). The anchor points are re-fitted to this strict form directly by ordinary least-squares regression, with no independent curvature parameter (Figure 6.1). With its curvature locked to the asymptote locations, the rigid form is constrained more strongly by points lying near the asymptotes than by points in the knee, and consequently returns sharper, more clearly defined asymptote estimates: $P_0 \approx 3.05$ kPa and $I_0 \approx 54.4$ kPa·ms. The pressure asymptote now sits at the upper edge of the probed reusable data rather than below it, the natural position for an asymptote inferred from data lying above it.

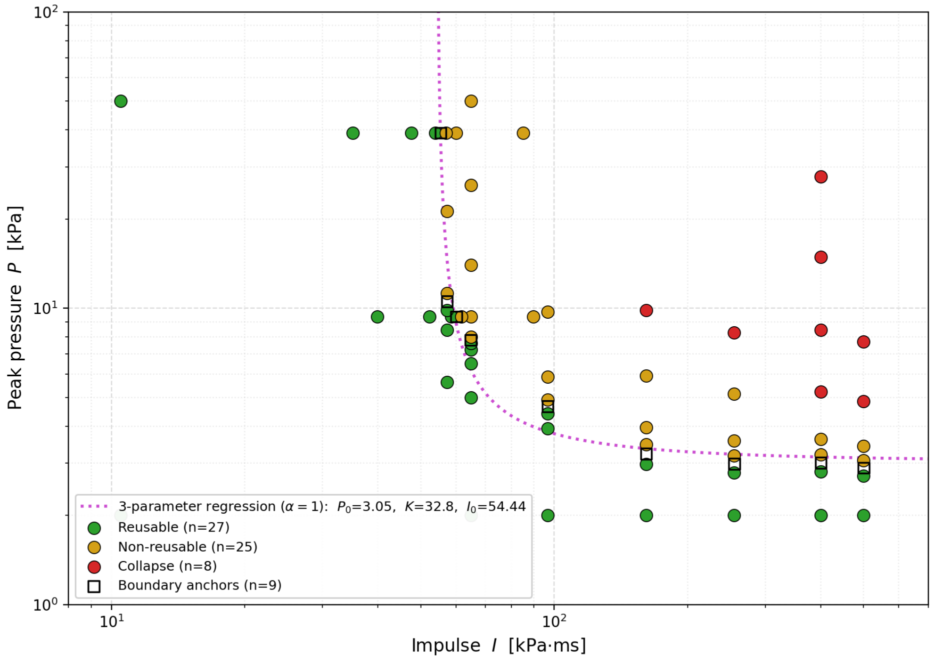

Figure 6.1: 3 parameter regression model on found anchor points

The disagreement between the three-parameter and four-parameter pressure asymptotes is informative. Both fits are consistent with the same anchor points; both are statistically defensible; and yet they bracket the physical pressure asymptote between 1.33 and 3.05 kPa. This range is the most honest characterisation of the asymptote available from the present sweep. A point estimate from either fit alone would over-state the precision of the result. The corresponding range for the impulse asymptote is narrower (the four-parameter fit returns $I_0 \approx 55.4$ kPa·ms against the three-parameter 54.4 kPa·ms), because the impulse asymptote is bracketed by data on both sides and is not dependent on extrapolation.

The price the three-parameter form pays for its tighter asymptotes is a less faithful representation of the dynamic transition. The three-parameter curve passes through the knee with a curvature dictated entirely by the asymptote locations and the scale parameter, with no degree of freedom available to match the local shape of the data. Where the four-parameter fit is anchor-data-rich in the knee and asymptote-extrapolating, the three-parameter fit is anchor-data-rich at the asymptotes and knee-extrapolating. The two fits do not contradict each other; they emphasise complementary aspects of the same boundary.

Piecewise interpolation as a check on the parametric form.

The second analysis abandons the parametric form altogether. Adjacent anchor points are connected by local exponential segments, producing a piecewise interpolation that reproduces every probe exactly (Figure 6.2). No global functional shape is imposed; the curve follows the data without smoothing.

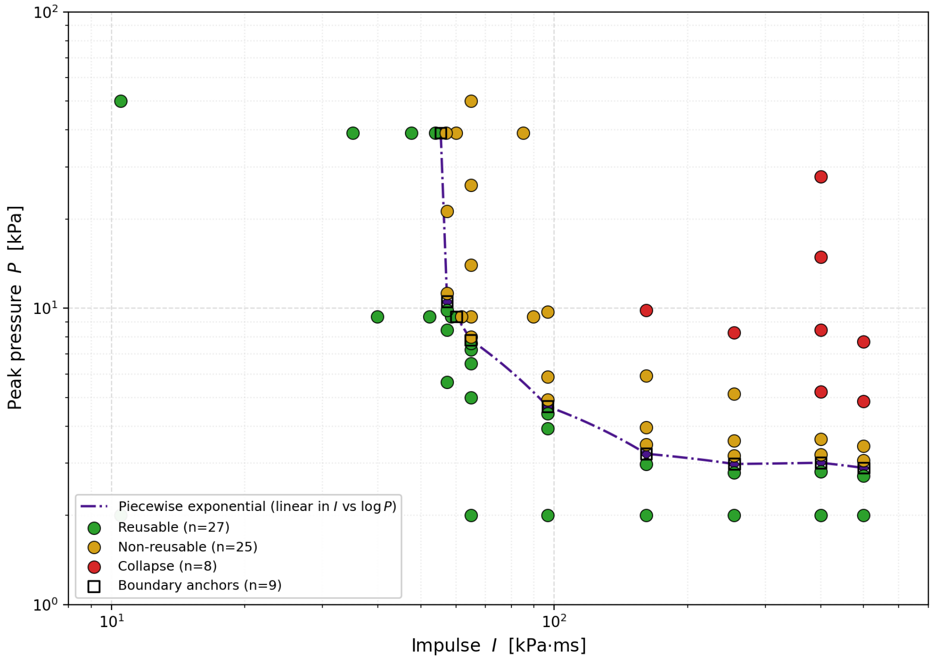

Figure 6.2: Piecewise exponential interpolation on found anchor points

The strength of this representation is the absence of any structural assumption about the underlying boundary. A wall whose true failure curve is not a rectangular hyperbola — a possibility that cannot be excluded *a priori* given the heterogeneity of materials and the variety of contact states involved — is represented more faithfully by a curve that follows the data than by one that smooths it. The piecewise curve in this case study exhibits no large kinks or discontinuities, which is consistent with a single failure mode governing the response across the probed range, but minor oscillations are visible between adjacent anchors. These oscillations are not large enough to suggest a mode transition; their origin is taken up in the next section.

The weakness of the piecewise representation is the reverse of its strength. The absence of a global functional shape means that every local irregularity in the anchor data, including those arising from numerical noise or from the finite per-bisection tolerance, is reproduced unfiltered. The curve cannot be extrapolated beyond the impulse range of the anchor points, and it cannot be summarised by a small set of scalar parameters. As a description of the data the piecewise representation is the most accurate of the three; as a description of the boundary it is the least compact.

6.3 Resolution, tolerance, and the case for more boundary points

The resolution of the reported boundary is set principally by the per-bisection tolerance: the width to which each bisection brackets the reusable–non-reusable transition along its axis before the anchor point is accepted. This is distinct from the number of anchor points and from where the sweep stops. The point count fixes how many points are placed and, with the zone-balancing of Section 3.3.4, how they are distributed along the boundary, but it does not change how tightly any single point is located; that is the tolerance alone.

For the present sweep the per-bisection tolerance on the boundary pressure was held at 0.5 kPa, and the sweep terminated on reaching the requested nine anchor points, with the evaluation budget left unconsumed. The tolerance does not influence the reported four-parameter fit directly, since the fit smooths over the uncertainty in individual anchors, but it sets the floor on how tightly each underlying point can be located and is therefore visible in the piecewise interpolation, which reproduces every anchor without smoothing. Between two adjacent anchors the piecewise curve cannot resolve features at a pressure scale finer than the tolerance, and at the lower end of the quasi-static asymptote, where the boundary itself lies near 3 kPa, the 0.5 kPa tolerance is a substantial fraction of the local pressure.

This is the principal explanation for the local oscillation visible in the piecewise interpolation. Several alternative explanations remain possible: physical multimodality of the response, numerical noise in the explicit solver, and contact-state sensitivity near the rotation threshold. None can be ruled out, but only the tolerance setting produces a change in the boundary that scales in proportion to the observed wobble. A finer per-bisection tolerance, at the cost of additional bisection iterations per anchor point, would compress the oscillation. The thesis sweep traded resolution for breadth, and the choice was governed less by the algorithm than by the wall-clock budget for the case study; with the budget available, the resolution chosen was appropriate to the question being asked, but the same configuration scaled to a finer tolerance would consume proportionally more computation. The same observation applies to the smoothed four-parameter fit indirectly: each anchor it is fitted against carries the same per-bisection uncertainty, and the curvature parameter n is the parameter most sensitive to small displacements of anchors in the knee region.

The number of anchor points itself is a separate axis. Nine anchors are sufficient to identify the qualitative shape of the boundary and the approximate location of both asymptotes, but they are sparse in the knee region (Figure 6.1) where the curvature is highest and where the four-parameter fit is most sensitive to small changes in the underlying data. Doubling the anchor budget to roughly twenty points, concentrated in the knee, would in principle resolve the curvature exponent more reliably and reduce the disagreement between the three-parameter and four-parameter asymptote estimates. Whether this is worthwhile depends on the application: for a P–I diagram intended for design tabulation, an asymptote uncertainty of a factor of two between fits is large; for a diagram intended only to indicate the order of magnitude of the failure threshold, it is acceptable.

A third axis, not exercised in the present case study, is the parameter-space coverage. The sweep ran on a single wall profile; the generality of the framework is asserted by the structure of the tool but not demonstrated by varying the wall geometry, the bond pattern, or the boundary conditions. The case study should therefore be read as a worked example of the tool's operation on a representative configuration rather than as a parametric study of URM blast response, which is a separate investigation.

6.4 Coverage relative to the ConWep-achievable domain

A useful perspective on the reported boundary is obtained by overlaying the final P–I curve on the ConWep-achievable domain mapped in Section 3.2.2 (Figure 6.3).

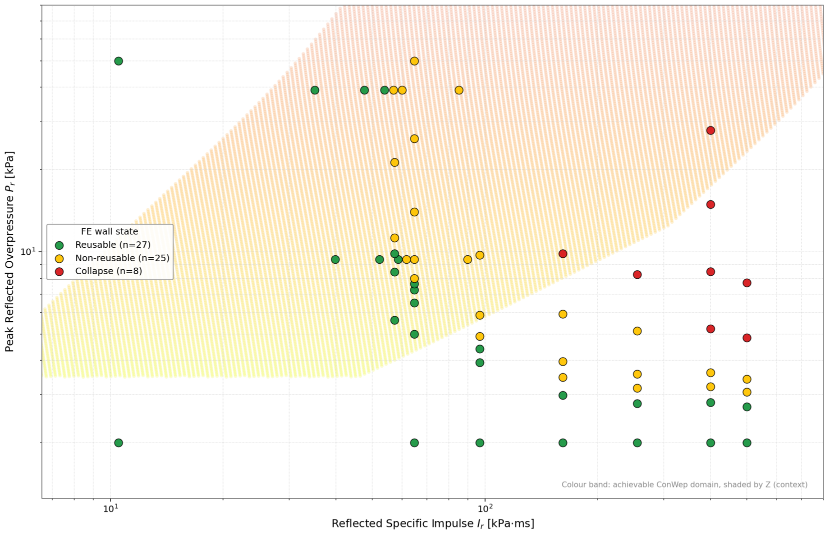

Figure 6.3: All evaluated points overlayed on achievable P-I domain through ConWep.

The impulsive branch of the boundary, in the vicinity of $I \approx 55$ to 65 kPa·ms and P on the order of tens of kilopascals, lies inside the achievable domain: a sufficiently small charge at a sufficiently small standoff produces a real-world ConWep pulse that probes this region. An analyst seeking to characterise the impulsive asymptote of a similar wall using a physically-grounded TNT-equivalent load model would therefore be able to do so within the validated regime of the Kingery–Bulmash equations.

The quasi-static branch of the boundary tells a different story. The pressure asymptote at $P_0 \approx 3$ kPa extends in principle to arbitrarily large impulses; the sweep probed it up to $I = 500$ kPa·ms. The combination of low overpressure and long impulse implied by this region is largely absent from the ConWep-achievable domain, because the real explosions catalogued by Kingery and Bulmash do not produce sustained low-pressure positive phases of this duration. An analyst restricted to ConWep loads would therefore be unable to characterise the quasi-static asymptote of this wall directly, and would necessarily fall back on extrapolation from the limited reachable region or on simplified analytical estimates.

This is the empirical complement to the argument made in Section 3.2.2 for adopting the single-parameter exponential pulse in sweep mode. The case study confirms, on a real boundary rather than on a theoretical coverage argument, that the failure boundary of a slender URM wall extends into a portion of the P–I plane that ConWep cannot reach, and that any sweep methodology limited to physically-realisable explosion pulses would miss a substantial part of the curve. The exponential pulse formulation adopted here removes this restriction at the cost of releasing the load from any direct correspondence with a particular explosion event — a cost that is acceptable in sweep mode, where the question being asked is parametric rather than scenario-specific.

6.5 Effectiveness of the flexural-mode isolation

The case-study wall was designed (Section 5.1) to isolate the flexural failure mode by combining a slender configuration ($H/t = 24$) with pinned out-of-plane and fixed in-plane support conditions. The intent of this design was to push the flexural P–I boundary downward in the P–I plane while pushing the boundaries of competing modes (sliding shear, in-plane shear, joint rocking) upward, so that the curve traced by the sweep would correspond unambiguously to one-way bending response.

The smooth, monotone shape of the fitted boundary is consistent with this picture but is not by itself sufficient evidence of mode isolation: an idealised hyperbolic fit will produce a smooth boundary even through anchor data that, in reality, straddles a mode transition, as long as the transition does not produce a discontinuity larger than the local fit residual.

A caveat must be entered nonetheless. The case study probed pressures up to 50 kPa and impulses up to 500 kPa·ms. A sufficiently impulsive load outside this range — the regime of contact or near-contact detonation, well above any (P, I) target considered here — could excite higher-order failure modes that the chosen configuration does not entirely suppress, and the mode-isolating wall profile is no guarantee of mode purity outside the probed window. The flexural attribution applies to the boundary in the parameter range reported, not to the wall's response under arbitrary loading.

6.6 Implications for the framework

Three implications for the framework follow from the case study.

First, the case study suggests that the reporting of the boundary may benefit from being decoupled from the parametric fit used internally by the algorithm. The current implementation produces and reports the four-parameter algorithmic fit as the single user-facing result, and the analysis of Section 6.2 indicates that, at least for this configuration, a single representation can be misleading: the asymptote it reports is an extrapolation, and the smooth global form it imposes obscures features that the located points themselves make visible. One alternative, motivated by this observation rather than established by it, would be to expose the anchor points as a first-class result, the piecewise interpolation as the default visualisation, and the three-parameter regression as a compact summary of the asymptote locations, with the four-parameter fit retained as an analytical expression for downstream design use but accompanied by an explicit caveat about its asymptote. Such a change would not affect the sweep algorithm, since the parametric fit remains essential during acquisition, only how the final output is constructed and presented. Whether it is warranted in general, rather than for the single boundary examined here, would need to be confirmed across a wider range of wall configurations.

Second, the resolution of the boundary is governed by the per-bisection tolerance, yet this quantity is not exposed to the user with the prominence its role warrants. The tolerance applies to every bisection in the sweep and therefore sets the floor on how tightly each anchor is located across the entire boundary. Its absolute value is fixed, so its relative effect grows as the boundary pressure falls: a fixed pressure tolerance is a negligible fraction of an anchor lying at tens of kilopascals but a large fraction of one lying near a few kilopascals, so the same setting over-resolves the high-pressure region and under-resolves the low-pressure tail. A future iteration of the GUI would benefit from exposing the per-bisection pressure tolerance as a direct input, ideally with a default tied to the local pressure scale so that the resolution is held roughly uniform in relative terms along the whole boundary rather than fixed in absolute terms.

Third, the framework's generality is demonstrated by its structure but not by the case study. A single wall configuration produces a single boundary, and the parametric study required to demonstrate that the boundary varies in physically reasonable ways with the wall geometry, slenderness, bond pattern, boundary conditions, and pre-stress, has not been performed. The structure of the GUI and the persistence of named profiles make such a parametric study mechanically straightforward; the cost is in wall-clock time. A short follow-up study running the same sweep configuration against three or four contrasting wall profiles, even at coarse anchor budgets, would substantially strengthen the claim that the tool generalises beyond the case-study wall.

7 Conclusion

This thesis set out to lower the analyst-time barrier to parametric blast assessment of unreinforced masonry walls by automating the construction of pressure–impulse failure boundaries. The work delivered a graphics-assisted tool, built on ABAQUS and driven through a parametric GUI, that composes three modelling choices into a single workflow that produces a P–I diagram without case-by-case manual model construction. The three choices are a simplified micro-model of the wall, with deformable brick units governed by Concrete Damaged Plasticity with rate-dependent hardening and mortar represented as zero-thickness cohesive interfaces; a dual-mode blast load, with ConWep in validation mode and a single-parameter exponential pulse in sweep mode; and an adaptive Bayesian boundary tracer. In this operational sense the design goal of the project was met: the user interacts only with dialog boxes and terminal commands, and the tool returns a fitted failure boundary together with the concrete boundary points on which it rests.

The framework was exercised end-to-end on a single-wythe, slender ($H/t = 24$) URM wall deliberately scoped to isolate out-of-plane flexure. The adaptive sweep located nine anchor points straddling the reusable–non-reusable transition in sixty finite-element evaluations, spanning the impulsive, dynamic, and quasi-static regimes, and terminated on the requested point count before the evaluation budget was exhausted. The resulting boundary exhibits the qualitative two-asymptote shape predicted by simplified single-degree-of-freedom theory, with an impulse asymptote near 54–$55\ \mathrm{kPa} \cdot \mathrm{ms}$ and a pressure asymptote bracketed between roughly 1.33 and $3.05\ \mathrm{kPa}$ depending on the fitted form. The smooth, monotone form of the boundary, with no kinks indicative of a competing failure mode crossing the dynamic regime, is consistent with the flexural response the wall configuration was designed to isolate, although this attribution rests on the design intent and the global boundary behaviour rather than on direct confirmation at the material level.

These results are presented as a demonstration of the pipeline's operation rather than a validated characterisation of a specific wall configuration. Several limitations bound the strength of the conclusion. First, the reported boundary is non-unique. The four-parameter fit the tool reports, a three-parameter regression, and a piecewise interpolation of the same anchor points yield curves that differ in the dynamic region and disagree on the location of the pressure asymptote. A single fitted curve therefore overstates the precision with which the asymptote is known. The quasi-static asymptote could in principle be bracketed independently by an elementary static-capacity calculation of the wall, which is not pursued here. Second, the boundary is internally consistent but externally unvalidated: the case-study wall is a fabricated configuration rather than a documented experimental specimen, so the absolute asymptote values cannot be checked against a benchmark; physical blast tests themselves carry large experimental scatter, so even a measured boundary would be an imprecise reference, while the fully specified case study presented here constitutes a reproducible deterministic benchmark against which future models and tools can be compared. Third, each evaluation simulates a finite response window, which in the quasi-static regime may end before peak midspan displacement is reached; where this occurs the support rotation is under-measured and the affected boundary points displaced, so the limitation can bias the low-pressure tail of the diagram in particular. Finally, the generality of the framework is asserted by its structure but not demonstrated, since only one wall configuration was swept: running the same procedure across contrasting profiles, varying slenderness, bond pattern, boundary conditions, and pre-stress, would show that the traced boundary shifts in physically reasonable ways with those parameters, which is the evidence that would substantiate the generality claim.

In conclusion, the methodology developed in this thesis establishes a reproducible route from a wall configuration to a traced failure boundary.

8 Future work

Several extensions to the work reported here are natural and would address limitations identified in Chapters 5 and 6.

8.1 Extension of the blast model

The single-wall scope adopted in this thesis allows the simplified ConWep load model to be used in validation mode without loss of accuracy, because no reflective or shadowing surfaces are present. Extending the model to include adjacent transverse walls, floors, or surrounding terrain would re-introduce the phenomena that distinguish CEL from ConWep, and adding CEL as a third selectable load model would become a natural extension.

8.2 Solver switching in the sweep

The explicit solver used throughout the sweep is well-suited to the impulsive and dynamic regions of the P–I plane but is wasteful in the quasi-static asymptote, where the transient response is slow and the loading is essentially static. Each evaluation in this region is correspondingly expensive in wall-clock time. A future iteration of the framework could detect when the target pulse falls in the quasi-static regime, potentially by monitoring an unusually long-running evaluation, and switch to an implicit solver for those cases, retaining the explicit solver only where it is needed.

8.3 Algorithmic improvements

The Bayesian adaptive sampling implemented here uses the predictive standard deviation as the acquisition function and the BIC criterion for parametric model promotion. A more efficient sampler could place evaluations more economically, concentrating effort in the regions where the current fit is poorest rather than spreading it evenly along the boundary. A more radical alternative is to abandon the parametric hyperbolic fit as the reported boundary entirely and to use the fit only as a guide for probe placement during sampling, reporting the boundary as the anchor points themselves with piecewise interpolation between them. This would let the user control the resolution of the final curve directly, at the cost of a stronger reliance on a tight per-bisection tolerance, since an interpolated curve reproduces every anchor without smoothing.

8.4 Local damage extraction and visualisation

The boundary reported in this thesis rests entirely on the global support-rotation criterion as currently proposed in the UFC [43]; local damage at the joint and brick level was not extracted, and the damage classification draws on no material-level information. Recovering and using local indicators is a natural line of further work.

A first step is to make local damage output meaningful, which the present modelling choices currently prevent in three ways. First, the cohesive contact quantities describing mortar opening and slip, and the brick-level plastic strain, are not recorded over the sweep and would have to be. Second, the known incompatibility in Abaqus between rate-dependent hardening and the damage evolution tables forces the Concrete Damaged Plasticity damage variables to zero, and would have to be resolved before those variables carry information. Third, the predominantly flexural regime that governs the reusable boundary does not drive the joints to separate and slide, so the wall would have to be loaded into ranges where they do before the local indicators register meaningful values.

Once a usable local signal exists, spatialising it is the natural complement to the global classification. A heat-map of accumulated mortar slip, crack opening, or brick plastic strain across the wall surface would show where damage concentrates rather than only whether the wall as a whole has passed a rotation threshold, which is information of direct interest to practitioners assessing retrofit priorities.

8.5 Parametric study of wall configurations

The case study reported in this thesis exercises a single wall profile, which shows that the tool runs but not that it produces sensible results for walls other than the one tested. Running the same sweep configuration against three or four contrasting profiles, varying slenderness, bond pattern, boundary conditions, and pre-stress, would show whether the traced boundary moves in physically expected ways, for example shifting downward as the wall is made more slender. Confirming that the boundary responds sensibly to these changes is what would demonstrate the generality of the framework, rather than leaving it asserted from the structure of the tool alone. The persistence of named profiles in the GUI makes such a study mechanically straightforward; the cost is in wall-clock time.

8.6 Validation of the tool against physical tests

The work reported here establishes a methodology and a working implementation, but it stops short of demonstrating that the boundaries the tool produces are quantitatively correct. As set out in Section 6.1, the case-study boundary was checked only for qualitative plausibility and internal consistency; it was not compared against a measured response, because the case-study wall is a fabricated configuration rather than a documented specimen. Turning the tool from a methodology into a predictive instrument would require validation against physical tests, of increasing cost and increasing evidential weight.

The most direct step is comparison against a documented blast experiment. The validation mode of the tool, in which a ConWep load reproduces a known charge weight and standoff, exists precisely for this purpose but was not exercised against a physical test in this thesis. Running the tool in validation mode against an instrumented experiment on a wall of comparable construction, and comparing the predicted support rotation, deflection history, and damage classification against the measured response, would provide the first direct evidence that the modelling choices reproduce real behaviour. A close match would lend confidence to boundaries swept for similar walls; a mismatch would localise which part of the chain, the constitutive model, the cohesive interface, the rate dependence, or the load representation, is responsible.

A more demanding form of validation is the calibration of individual modelling choices against dedicated material and component tests. Several inputs to the present model are adopted from the literature or assumed in the absence of data, the tensile rate dependence taken equal to the compressive Dynamic Increase Factor being the clearest example. Targeted tests, of brick and mortar properties at the relevant strain rates, of cohesive interface behaviour, and of single-wall response under controlled dynamic loading, would replace these assumptions with measured values specific to the masonry being assessed. This is the point at which the tool would cease to rely on generic literature parameters and become capable of configuration-specific prediction. Together with the parametric study of Section 8.5, which tests whether the tool behaves sensibly across its intended domain, these tests would constitute the evidence required to use the tool's boundaries for protective assessment with quantitative rather than indicative confidence.

9 References

[1] S. M. Anas, M. Alam, and M. Umair. "Strengthening of braced unreinforced brick masonry wall with (i) C-FRP wrapping, and (ii) steel angle-strip system under blast loading". In: *Materials Today: Proceedings* 58 (2022), pp. 1181–1198.

[2] W. E. Baker et al. *Explosion Hazards and Evaluation*. Vol. 5. Fundamental Studies in Engineering. Amsterdam: Elsevier Scientific Publishing Company, 1983.

[3] David J. Benson. "Computational methods in Lagrangian and Eulerian hydrocodes". In: *Computer Methods in Applied Mechanics and Engineering* 99.2–3 (1992), pp. 235–394.

[4] Barron J. Bichon et al. "Efficient global reliability analysis for nonlinear implicit performance functions". In: *AIAA Journal* 46.10 (2008), pp. 2459–2468.

[5] J. E. Bolander et al. "Rigid-Body-Spring Network modeling of cement-based composites". In: *Fracture Mechanics of Concrete Structures*. Lisse: Swets & Zeitlinger, 2001, pp. 417–424.

[6] I. Bouckaert, Michele Godio, and J. P. de Almeida. "Nonlinear Response-History Analyses of Masonry and Mixed Structures With HybriDFEM". In: *Earthquake Engineering & Structural Dynamics* 55.8 (2026), pp. 1694–1709. DOI: 10.1002/eqe.70127.

[7] Siro Casolo. "Macroscale modelling of the orthotropic shear damage in the dynamics of masonry towers by RBSM". In: *Engineering Failure Analysis* 130 (2021), p. 105744.

[8] Siro Casolo. "Modelling in-plane micro-structure of masonry walls by rigid elements". In: *International Journal of Solids and Structures* 41.13 (2004), pp. 3625–3641.

[9] Leonid Chernin, Meir Vilnay, and Demetris M. Cotsovos. "Extended P-I diagram method". In: *Engineering Structures* 224 (2020), p. 111217.

[10] Leonid Chernin et al. "Pressure-impulse diagram method – a fundamental review". In: *Proceedings of the Institution of Civil Engineers – Engineering and Computational Mechanics* 172.2 (2019), pp. 55–69.

[11] Dassault Systèmes Simulia Corp. *Abaqus Analysis User's Guide*. Providence, RI, USA, 2024.

[12] K. Doherty et al. "Displacement-based seismic analysis for out-of-plane bending of unreinforced masonry walls". In: *Earthquake Engineering & Structural Dynamics* 31.4 (2002), pp. 833–850.

[13] Benjamin Echard, Nicolas Gayton, and Maurice Lemaire. "AK-MCS: An active learning reliability method combining Kriging and Monte Carlo Simulation". In: *Structural Safety* 33.2 (2011), pp. 145–154.

[14] Daniel Foreman-Mackey et al. "emcee: The MCMC Hammer". In: *Publications of the Astronomical Society of the Pacific* 125.925 (2013), pp. 306–312.

[15] F. G. Friedlander. “The diffraction of sound pulses. I. Diffraction by a semi-infinite plane”. In: *Proceedings of the Royal Society of London A* 186.1006 (1946), pp. 322–344.

[16] Michele Godio and Mathias Flansbjer. *Blast capacity of masonry walls – Modelling of static and impact loads*. CFORT Work Report 2025:04. ISSN 3119-2726. Centrum för Fortifikatorisk Kompetens (CFORT), 2025.

[17] Michele Godio and Mathias Flansbjer. “Experimental response of unreinforced brick masonry walls to moderate-velocity point-load and line-load impacts”. In: *International Journal of Impact Engineering* 206 (2025), p. 105461. DOI: 10.1016/j.ijimpeng.2025.105461.

[18] Michele Godio and Mathias Flansbjer. “FEM meso-scale modelling of brick walls subjected to impacts and blasts: Formulation and laboratory test validation”. In: *Proceedings of the 15th International Conference on Shock & Impact Loads on Structures*. 2025, pp. 255–263. DOI: 10.63959/silos15.

[19] Michele Godio, Mathias Flansbjer, and Natalie Williams Portal. “Low-velocity out-of-plane impact tests on double-wythe unreinforced brick masonry walls instrumented with optical measurements”. In: *International Journal of Impact Engineering* 178 (2023), p. 104597. DOI: 10.1016/j.ijimpeng.2023.104597.

[20] Michele Godio, Mathias Flansbjer, and Natalie Williams Portal. “Single- and double-wythe brick masonry walls subjected to four-point bending tests under different support conditions: Simply supported, rigid, non-rigid”. In: *Construction and Building Materials* 404 (2023), p. 132544. DOI: 10.1016/j.conbuildmat.2023.132544.

[21] Michele Godio et al. “Experimental and numerical approaches to investigate the out-of-plane response of unreinforced masonry walls subjected to free far-field blasts”. In: *Engineering Structures* 239 (2021), p. 112328. DOI: 10.1016/j.engstruct.2021.112328.

[22] Jonathan Goodman and Jonathan Weare. “Ensemble samplers with affine invariance”. In: *Communications in Applied Mathematics and Computational Science* 5.1 (2010), pp. 65–80.

[23] S. S. Hameed, J. Jose, and P. Nagarajan. “Shear Strengthening of RC Beams Using Various Material Laminates”. In: *Advances in Structural Mechanics and Applications: Proceedings of ASMA-2021*. Springer Nature Switzerland, 2023, pp. 488–507.

[24] Hong Hao and Boris G. Tarasov. “Experimental study of dynamic material properties of clay brick and mortar at different strain rates”. In: *Australian Journal of Structural Engineering* 8.2 (2008), pp. 117–132.

[25] David W. Hyde. *Conventional Weapons Program (CONWEP): User’s Reference Manual*. U.S. Army Engineer Waterways Experiment Station. Vicksburg, MS, 1991.

[26] Donald R. Jones, Matthias Schonlau, and William J. Welch. “Efficient global optimization of expensive black-box functions”. In: *Journal of Global Optimization* 13.4 (1998), pp. 455–492.

[27] Robert E. Kass and Adrian E. Raftery. “Bayes factors”. In: *Journal of the American Statistical Association* 90.430 (1995), pp. 773–795.

[28] Tsuneo Kawai. “New discrete models and their application to seismic response analysis of structures”. In: *Nuclear Engineering and Design* 48.1 (1978), pp. 207–229.

[29] Charles N. Kingery and Gerald Bulmash. *Airblast Parameters from TNT Spherical Air Burst and Hemispherical Surface Burst*. Tech. rep. ARBRL-TR-02555. Aberdeen Proving Ground, Maryland: U.S. Army Armament Research and Development Center, Ballistic Research Laboratory, 1984.

[30] Theodor Krauthammer et al. “Pressure-impulse diagrams for the behavior assessment of structural components”. In: *International Journal of Impact Engineering* 35.8 (2008), pp. 771–783.

[31] E. L. Lee, H. C. Hornig, and J. W. Kury. *Adiabatic Expansion of High Explosive Detonation Products*. Tech. rep. TID-4500-UCRL-50422. Livermore, CA: Lawrence Radiation Laboratory, University of California, 1968.

[32] Paulo B. Lourenço. “Computations on historic masonry structures”. In: *Progress in Structural Engineering and Materials* 4.3 (2002), pp. 301–319.

[33] Johan Magnusson, Mikael Hallgren, and Anders Ansell. “Shear in concrete structures subjected to dynamic loads”. In: *Structural Concrete* 15.1 (2014), pp. 55–65.

[34] Fulvio Parisi, Claudio Balestrieri, and Domenico Asprone. “Blast resistance of tuff stone masonry walls”. In: *Engineering Structures* 113 (2016), pp. 233–244.

[35] Viktor Peterson. “Shear-type failure of concrete structural elements under blast and impact”. TRITA-ABE-DLT-2547. Doctoral Thesis in Civil and Architectural Engineering. Stockholm, Sweden: KTH Royal Institute of Technology, 2026.

[36] S. H. Rafsanjani, Paulo B. Lourenço, and N. Peixinho. “Dynamic interface model for masonry walls subjected to high strain rate out-of-plane loads”. In: *International Journal of Impact Engineering* 76 (2015), pp. 28–37.

[37] Gideon Schwarz. “Estimating the dimension of a model”. In: *The Annals of Statistics* 6.2 (1978), pp. 461–464.

[38] Bobak Shahriari et al. “Taking the human out of the loop: A review of Bayesian optimization”. In: *Proceedings of the IEEE* 104.1 (2016), pp. 148–175.

[39] S. Shamim, S. Ahmad, and R. A. Khan. “Dynamic Response of Masonry Panel Under High-Strain Loading Condition”. In: *Advances in Structural Mechanics and Applications: Proceedings of ASMA-2021*. Springer Nature Switzerland, 2023, pp. 78–89.

[40] G. Svedbjörk. *Elastic beam subjected to transient load: Moment and shear distributions in time and space*. Publication 42. Eskilstuna, Sweden: Swedish Fortifications Agency, 1975.

[41] Swedish Fortifications Agency. *Fortifikationsverkets konstruktionsregler FKR 2011 (Design Manual for Protective Structures)*. Dnr 4535/2011. Eskilstuna, Sweden, 2011.

[42] M. M. Swisdak. “Simplified Kingery airblast calculations”. In: *Proceedings of the 26th DoD Explosives Safety Seminar*. Miami, FL: U.S. Department of Defense Explosives Safety Board, 1994.

[43] United States Department of Defense. *Unified Facilities Criteria (UFC) 3-340-02: Structures to Resist the Effects of Accidental Explosions*. 5 December 2008. Washington, DC, Dec. 2008.

[44] Xing Wei, Ting Huang, and N. Li. “Numerical derivation of pressure-impulse diagrams for unreinforced brick masonry walls”. In: *Advanced Materials Research* 368–373 (2011), pp. 1435–1439.

[45] B. J. Zapata and D. C. Weggel. “Collapse study of an unreinforced masonry bearing wall building subjected to internal blast loading”. In: *Journal of Performance of Constructed Facilities* 22.2 (2008), pp. 92–100.

DEPARTMENT OF ARCHITECTURE AND
CIVIL ENGINEERING
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2026
www.chalmers.se

AVANCEZ
1829

CHALMERS
UNIVERSITY OF TECHNOLOGY